Lex Fridman

Lecture 5:

Deep Learning for Human Sensing

Deep Learning for Human Sensing

- Requirements for success (from more to less critical)
  - **Data:** A lot of real-world data (and algorithms that learn from data)
  - **Semi-supervised:** Human annotations of representative subsets of data
  - **Efficient annotation:** Specialized annotation tooling
  - **Hardware:** Large-scale distributed compute and storage
  - **Robustness:** Algorithms that don't need calibration (learn the calibration)
  - **Temporal dynamics:** Algorithms that consider time
- Current importance relation for successful application of deep learning:

>

Good Algorithms*

* As long as they learn from data

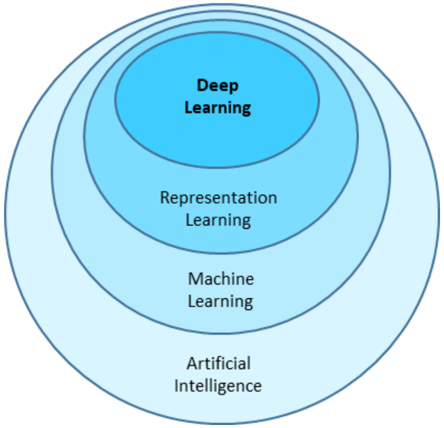

Overview

- **Human Imperfections**
- Pedestrian Detection
- Body Pose Estimation
- Glance Classification
- Emotion Recognition
- Cognitive Load Estimation
- Human-Centered Vision for Autonomous Vehicles

Humans Are Amazing

Humans Are Amazing

- 3.22 trillion miles (US, 2016)
- 40,200 fatalities (US, 2016)
- 1 fatality per 80 million miles
- 1 in 625 chance of dying in car crash (in your lifetime)

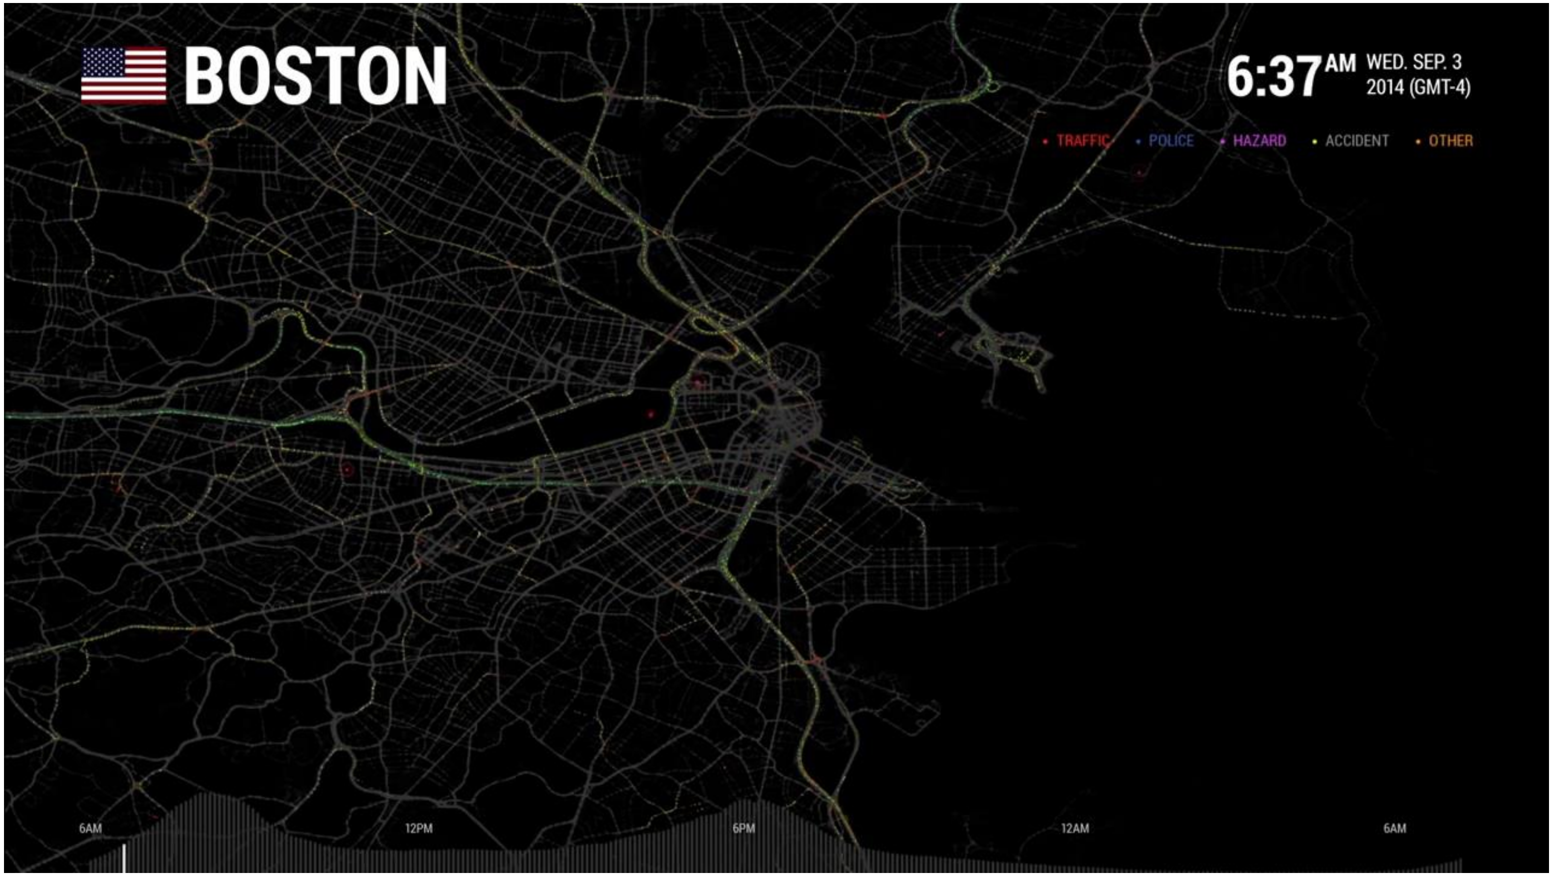

Humans are Flawed

What is distracted driving?

- Texting
- Using a smartphone
- Eating and drinking
- Talking to passengers
- Grooming
- Reading, including maps
- Using a navigation system
- Watching a video
- Adjusting a radio

- **Injuries and fatalities:**
  3,179 people were killed and 431,000 were injured in motor vehicle crashes involving distracted drivers
  (in 2014)
- **Texts:**
  169.3 billion text messages were sent in the US every month.
  (as of December 2014)
- **Eye off road:**
  5 seconds is the average time your eyes are off the road while texting. When traveling at 55mph, that's enough time to cover the length of a football field blindfolded.

Humans are Flawed

- **Drunk Driving:** In 2014, 31 percent of traffic fatalities involved a drunk driver.
- **Drugged Driving:** 23% of night-time drivers tested positive for illegal, prescription or over-the-counter medications.
- **Distracted Driving:** In 2014, 3,179 people (10 percent of overall traffic fatalities) were killed in crashes involving distracted drivers.
- **Drowsy Driving:** In 2014, nearly three percent of all traffic fatalities involved a drowsy driver, and at least 846 people were killed in crashes involving a drowsy driver.

Two Paths to an Autonomous Future

A1:
Human-Centered Autonomy

Blue Text: Easier

Red Text: Harder

- **Localization and Mapping:** Where am I?
- **Scene Understanding:** Where/who/what/why of everyone else?
- **Movement Planning:** How do I get from A to B?
- **Human-Robot Interaction:** What is the physical and mental state of the driver?
- **Communicate:** How to I convey intent to the driver and to the world?

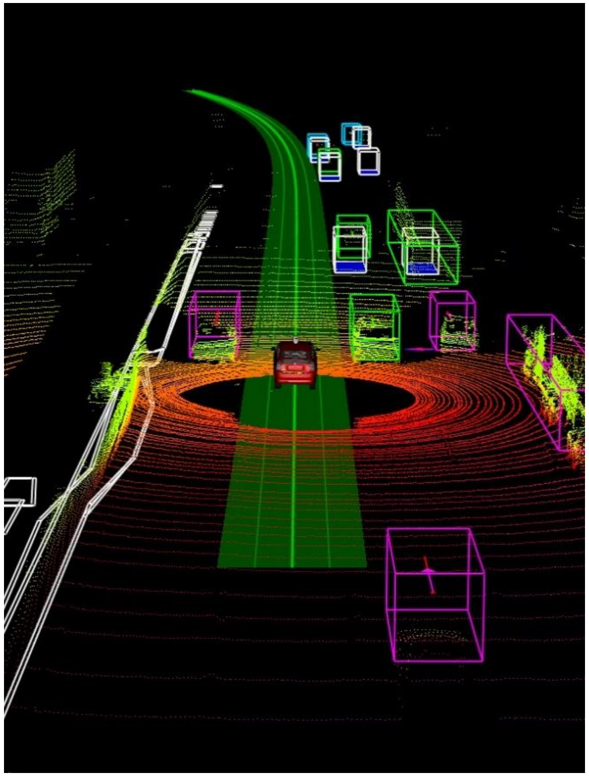

A2:
Full Autonomy

- **Localization and Mapping:** Where am I?
- **Scene Understanding:** Where/who/what/why of everyone else?
- **Movement Planning:** How do I get from A to B?
- **Human-Robot Interaction:** What is the physical and mental state of the driver?
- **Communicate:** How to I convey intent to the driver and to the world?

Is partially automated driving a bad idea? Observations from an on-road study

Article · April 2018 *with* 447 Reads
DOI: 10.1016/j.apergo.2017.11.010

Victoria Banks
14.44 · University of Southampton

Alexander Eriksson
11.13 · Swedish National Road and Transport Research Inst...

Jim O'donoghue

Neville A Stanton
43.23 · University of Southampton

WAYMO

Chris Urmson

Public Perception of What Drivers Do in Semi-Autonomous Vehicles

Public Perception of What Drivers Do in Semi-Autonomous Vehicles

MIT-AVT Naturalistic Driving Dataset

Vehicles instrumented: 25

Distance traveled: 275,000+ miles

Video frames: 4.7+ billion

Hardware

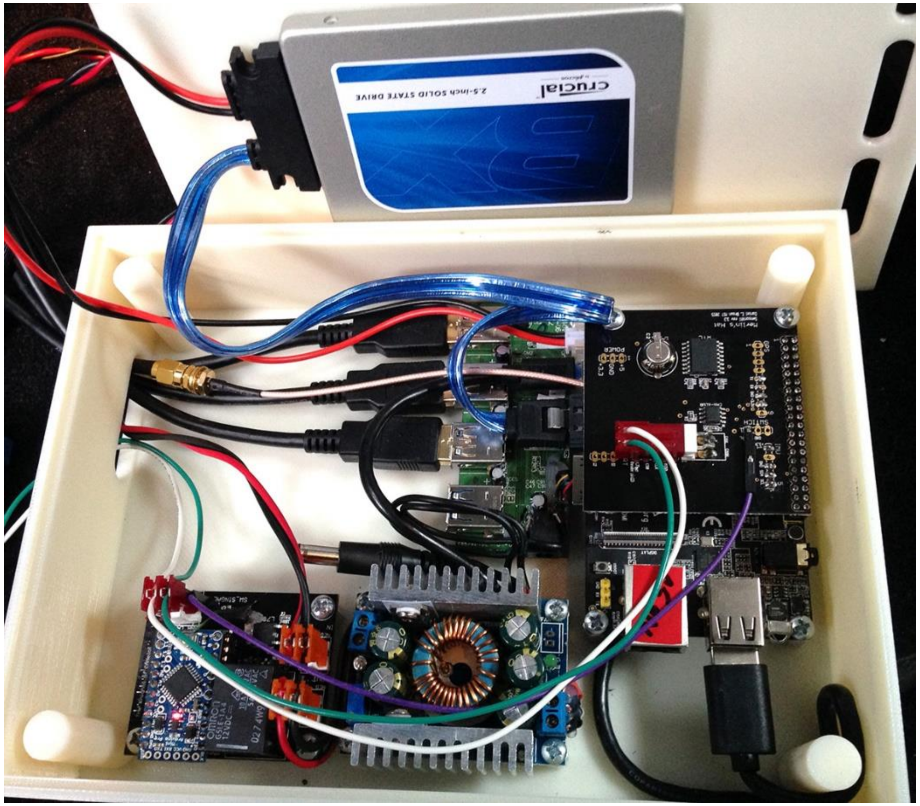

Software

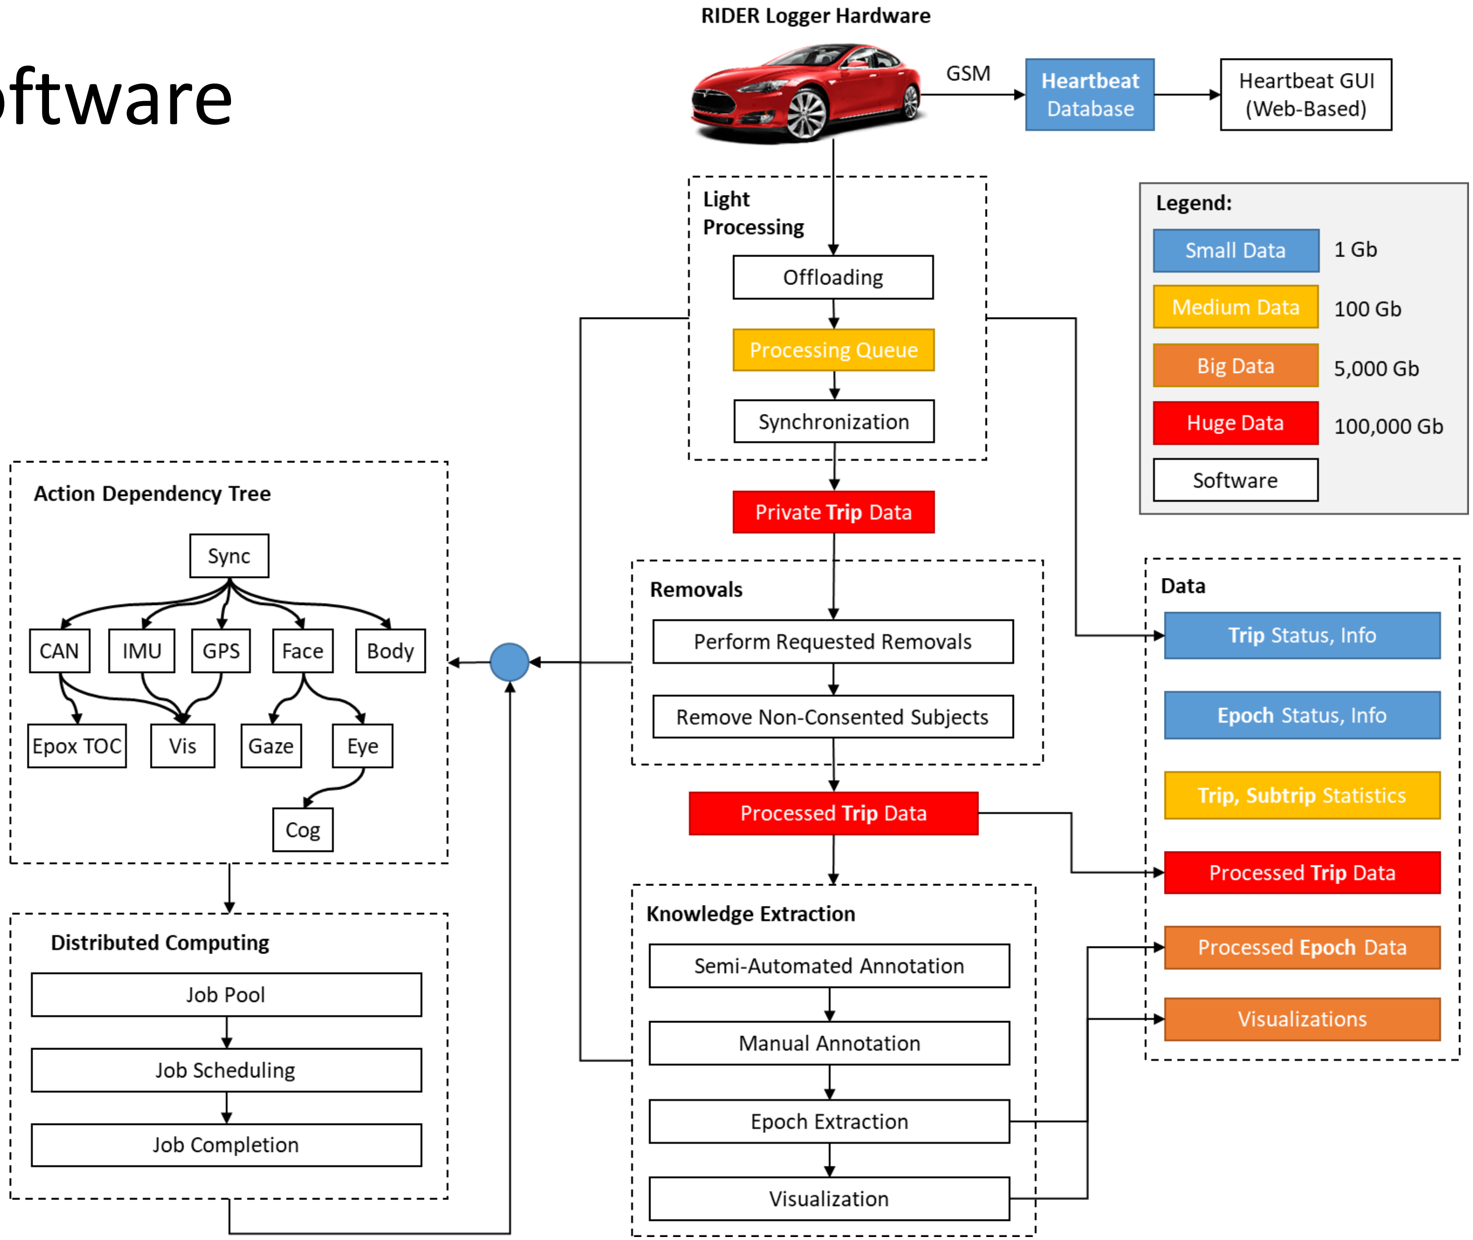
RIDER Logger Hardware
GSM
Heartbeat Database
Heartbeat GUI (Web-Based)
Legend:
Small Data
1 Gb
Medium Data
100 Gb
Big Data
5,000 Gb
Huge Data
100,000 Gb
Software
Light Processing
Offloading
Processing Queue
Synchronization
Private Trip Data
Removals
Perform Requested Removals
Remove Non-Consented Subjects
Processed Trip Data
Knowledge Extraction
Semi-Automated Annotation
Manual Annotation
Epoch Extraction
Visualization
Data
Trip Status, Info
Epoch Status, Info
Trip, Subtrip Statistics
Processed Trip Data
Processed Epoch Data
Visualizations
Action Dependency Tree
Sync
CAN
IMU
GPS
Face
Body
Epox TOC
Vis
Gaze
Eye
Cog
Distributed Computing
Job Pool
Job Scheduling
Job Completion


Total Time Driving: 0 mins
Autopilot Available: 0 mins
Autopilot Engaged: 0 mins

Human Behavior

Shared Autonomy

Understand Behavior

Assist Behavior

Share Control

Semi-Supervised Learning

Large-Scale Naturalistic Data

MIT-AVT Naturalistic Driving Dataset

MIT Autonomous Vehicle Technology Study

Study months to-date: **21**
Participant days: **7,146**
Drivers: **78**
Vehicles: **25**
Miles driven: **275,589**
Video frames: **3.48 billion**

Study data collection is ongoing.
*Statistics updated on: **Oct 23, 2017**.*

Tesla Model S
24,657 miles
588 days in study

Tesla Model X
22,001 miles
421 days in study

Tesla Model S
18,896 miles
435 days in study

Tesla Model S
18,666 miles
353 days in study

Range Rover Evoque
18,130 miles
483 days in study

Tesla Model S
15,735 miles
322 days in study

Tesla Model X
15,074 miles
276 days in study

Range Rover Evoque
14,499 miles
440 days in study

Tesla Model S
14,410 miles
371 days in study

Tesla Model S
14,117 miles
248 days in study

Volvo S90
13,970 miles
325 days in study

Tesla Model S
12,353 miles
321 days in study

Volvo S90
11,072 miles
412 days in study

Tesla Model X
10,271 miles
366 days in study

Tesla Model S
9,188 miles
183 days in study

Tesla Model S
8,319 miles
374 days in study

Tesla Model S
6,720 miles
194 days in study

Tesla Model S
5,186 miles
91 days in study

Tesla Model X
5,111 miles
232 days in study

Tesla Model S
4,596 miles
132 days in study

Tesla Model X
4,587 miles
233 days in study

Tesla Model X
3,719 miles
133 days in study

Tesla Model S
3,006 miles
144 days in study

Tesla Model X
1,306 miles
69 days in study

Tesla Model S
(Offload pending)

500+ Miles / Day and Growing

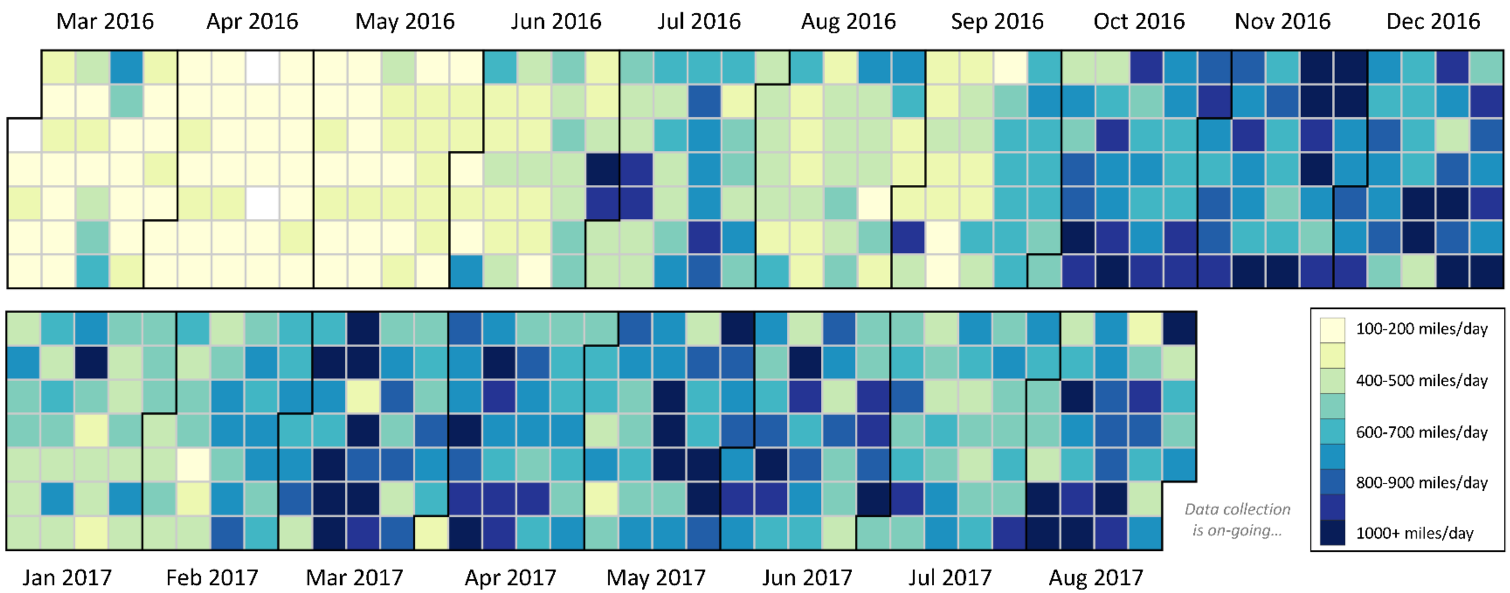

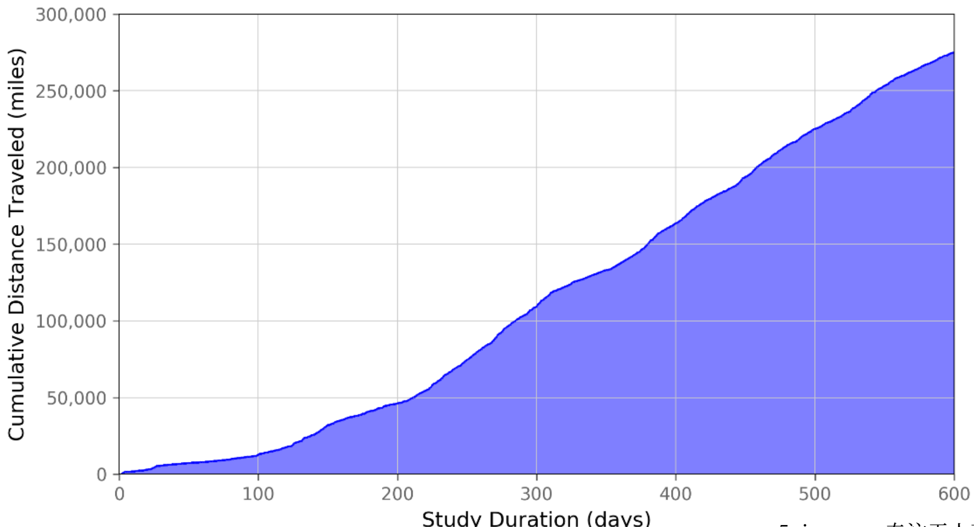

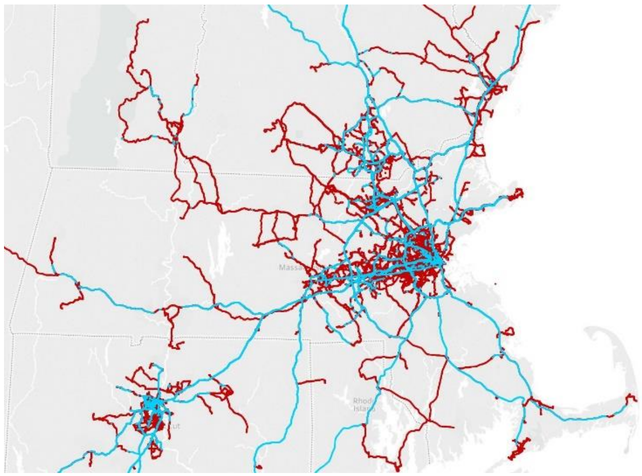

Tesla Autopilot: Patterns of Use

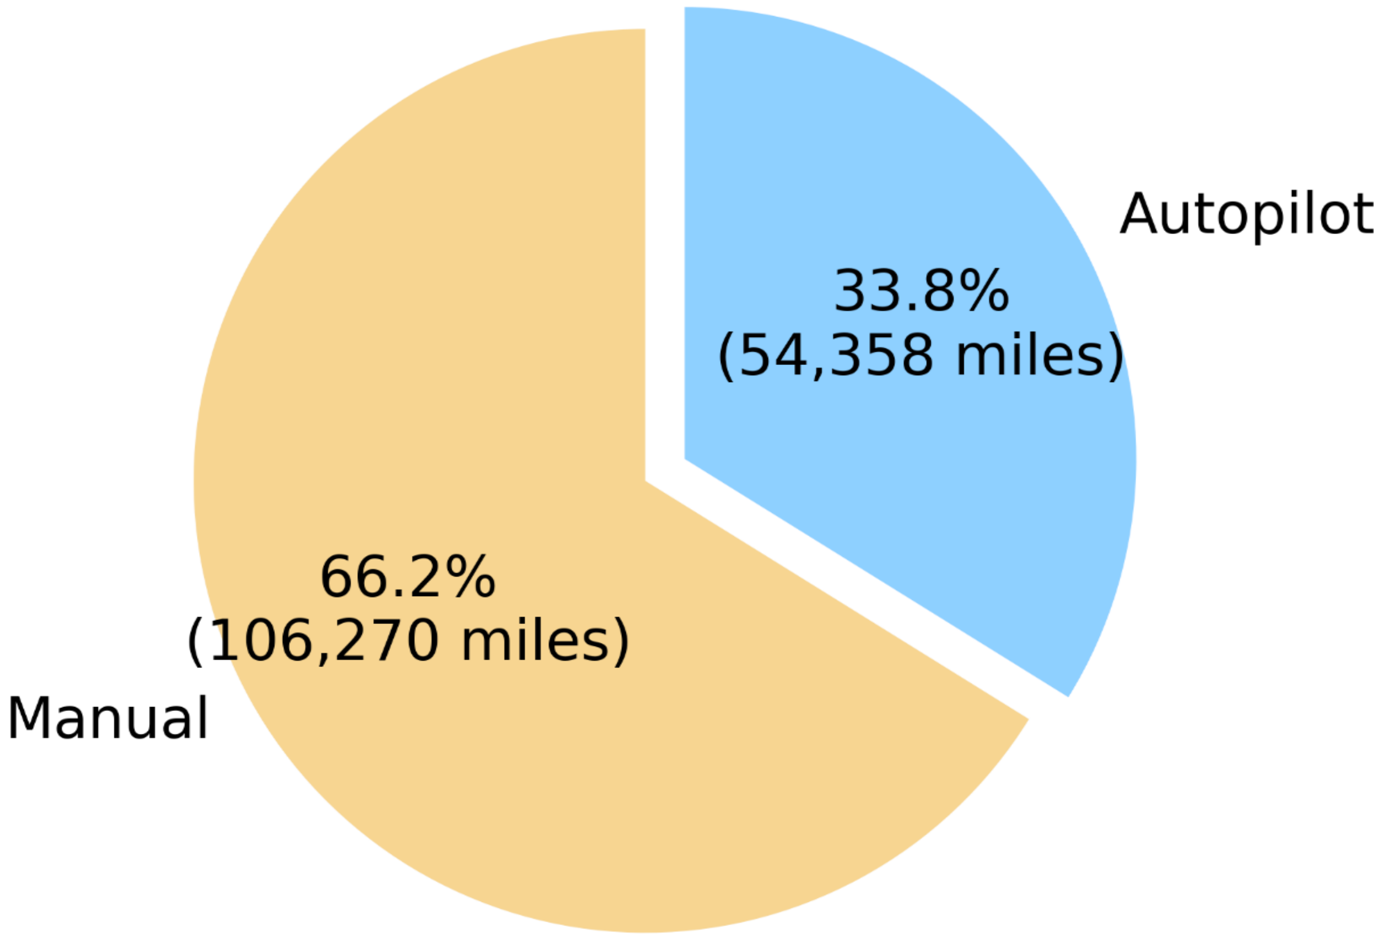

33.8% of the **miles** driven are with Autopilot engaged

Physical Engagement: Glance Classification

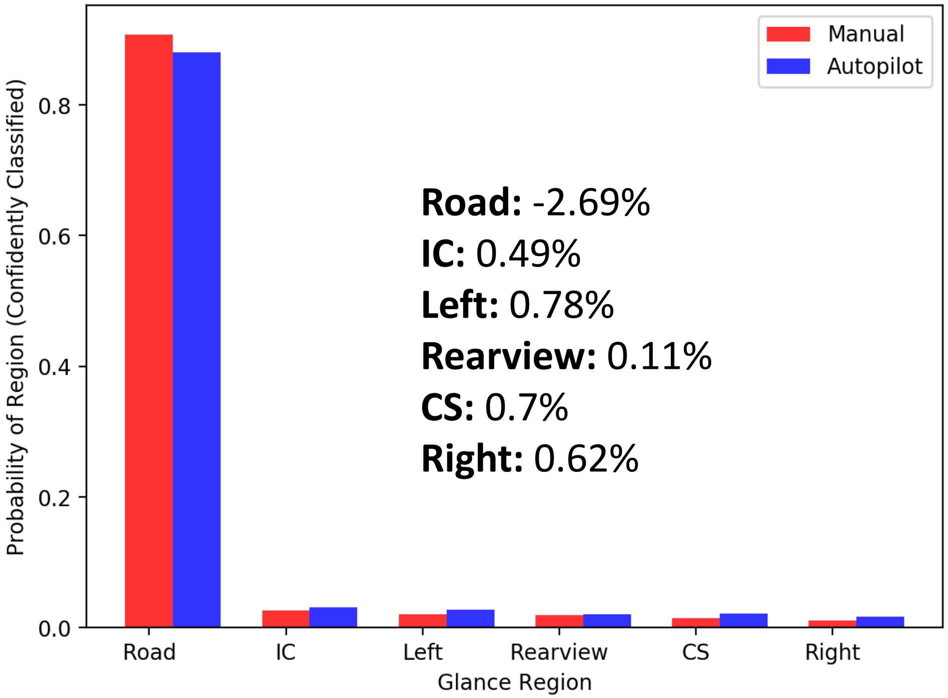

Semi-Autonomous Driving: Observed Patterns of Behavior

- The "how" of successful human-robot interaction:

  Use but Don't Trust.

- The "why" of successful human-robot interaction:

  Learn Limitations by Exploring.

Deep Learning for Human Sensing

- Requirements for success (from more to less critical)
  - **Data:** A lot of real-world data (and algorithms that learn from data)
  - **Semi-supervised:** Human annotations of representative subsets of data
  - **Efficient annotation:** Specialized annotation tooling
  - **Hardware:** Large-scale distributed compute and storage
  - **Robustness:** Algorithms that don’t need calibration (learn the calibration)
  - **Temporal dynamics:** Algorithms that consider time
- Current importance relation for successful application of deep learning:

>

Good Algorithms*

* As long as they learn from data

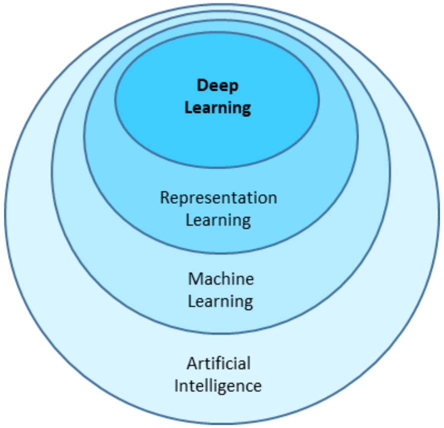

Overview

- Human Imperfections
- **Pedestrian Detection**
- Body Pose Estimation
- Glance Classification
- Emotion Recognition
- Cognitive Load Estimation
- Human-Centered Vision for Autonomous Vehicles

Human Sensing:
A Deep Learning Perspective

Increasing level of detection resolution and **difficulty**

Pedestrian Detection | Body Pose | Head Pose | Blink Rate | Blink Duration | Eye Pose | Blink Dynamics | Pupil Diameter | Micro Saccades

Face Detection | Face Classification | Glance Classification | Drowsiness | Micro Glances | Cognitive Load

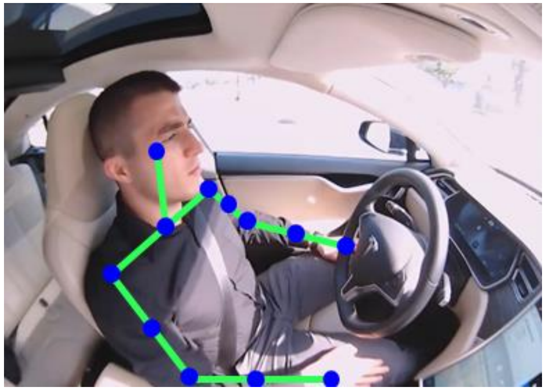

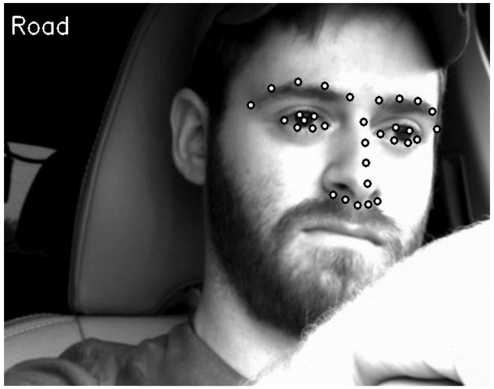

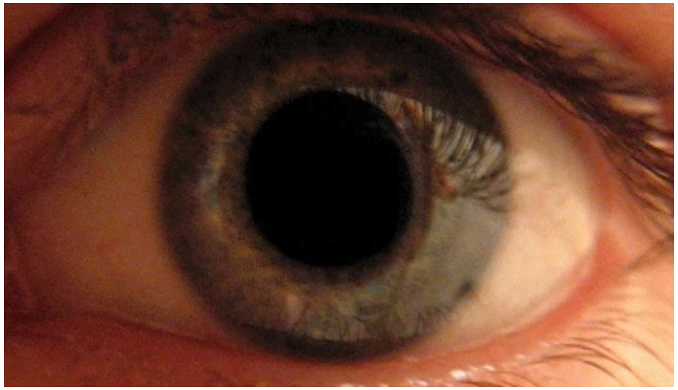

Pedestrian Detection

- The usual challenges, e.g.:
  - Various style of clothing in appearance
  - Different possible articulations
  - The presence of occluding accessories
  - Frequent occlusion between pedestrians

- History of object detection
  - Sliding window
    - Haar Cascades
    - Histogram of Oriented Features
    - CNN
  - R-CNN, Fast R-CNN, Faster R-CNN
  - Mask RCNN (adds segmentation)
  - VoxelNet (detection in 3D space)

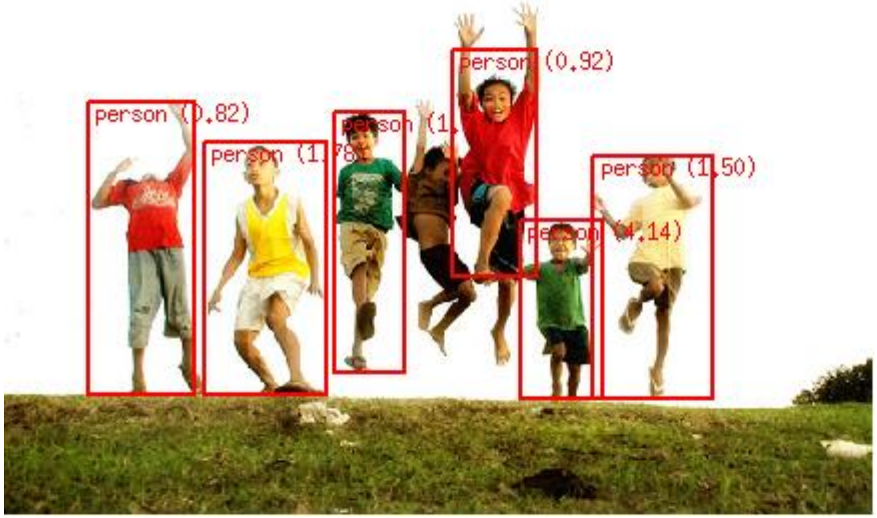

R-CNN: Regions with CNN Features

- Simple algorithm
  - Extract region proposals (selective search)
  - Use CNN on each one (w/ non-maximum suppression)

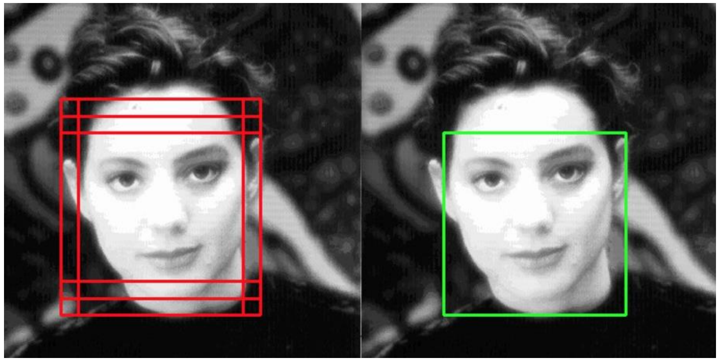

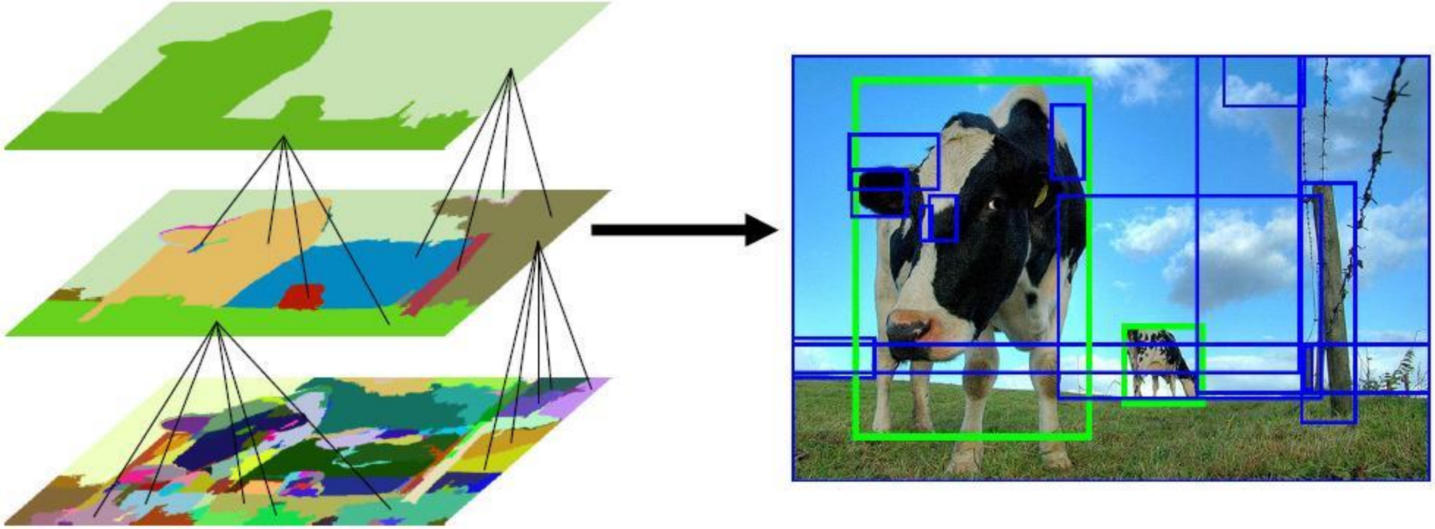

- Per 10 hours (1 recording day)
  - 12,000 pedestrians
  - 21,600,000 samples of feature vector

Naturalistic Driving Data:

Pedestrians, Cyclists, Other Cars

Sony FDR-AX53

ZED Stereo Camera

Gear 360 Camera

GoPro Hero4

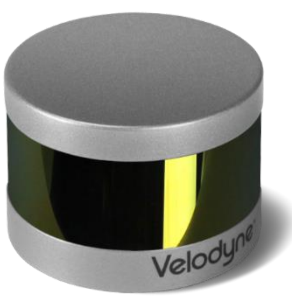

Velodyne VLP-16

Velodyne HDL-64E

Naturalistic Driving Data:

Pedestrians, Cyclists, Other Cars

Naturalistic Driving Data:

Pedestrians, Cyclists, Other Cars

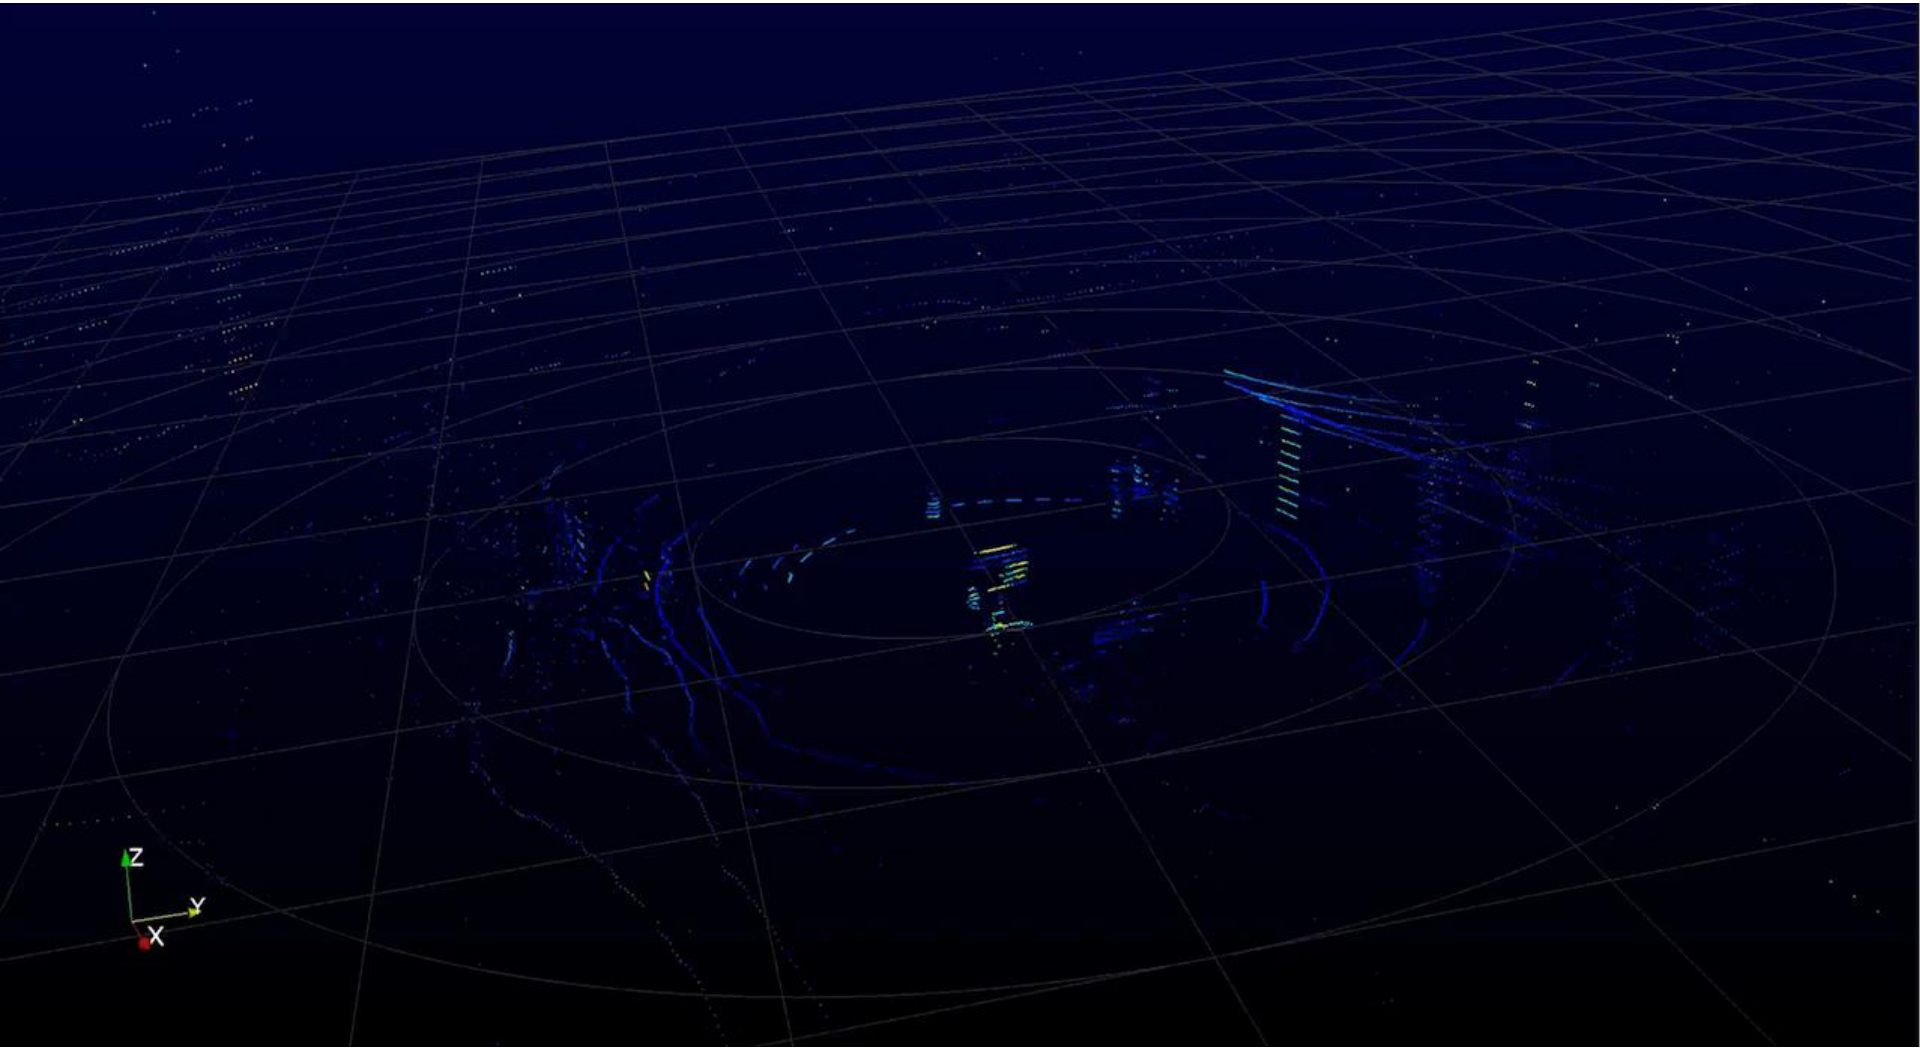

Naturalistic Driving Data:

Pedestrians, Cyclists, Other Cars

Naturalistic Driving Data:

Pedestrians, Cyclists, Other Cars

Pedestrian Detection

Overview

- Human Imperfections
- Pedestrian Detection
- **Body Pose Estimation**
- Glance Classification
- Emotion Recognition
- Cognitive Load Estimation
- Human-Centered Vision for Autonomous Vehicles

Human Sensing:
A Deep Learning Perspective

Increasing level of detection resolution and **difficulty**

Pedestrian Detection | Body Pose | Head Pose | Blink Rate | Blink Duration | Eye Pose | Blink Dynamics | Pupil Diameter | Micro Saccades

Face Detection | Face Classification | Glance Classification | Drowsiness | Micro Glances | Cognitive Load

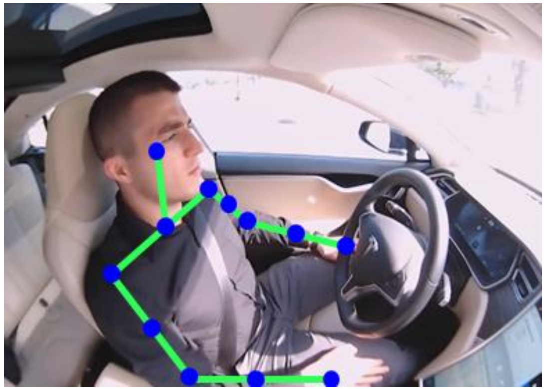

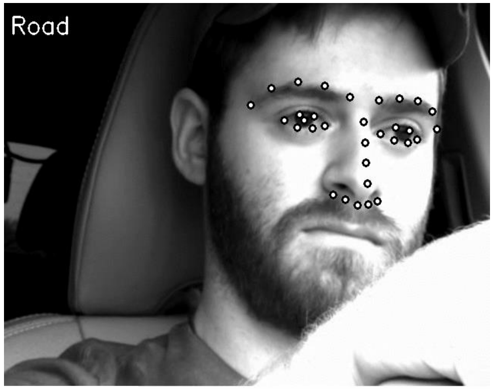

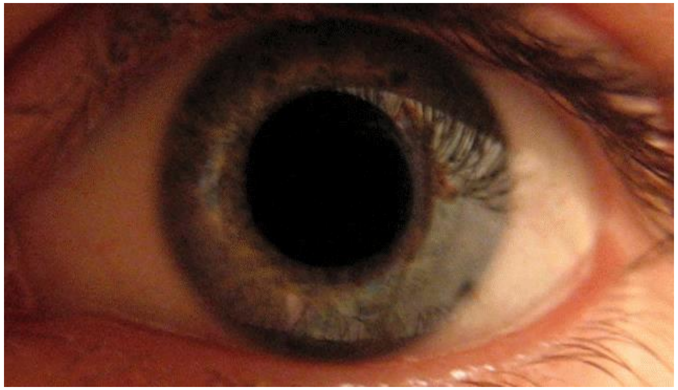

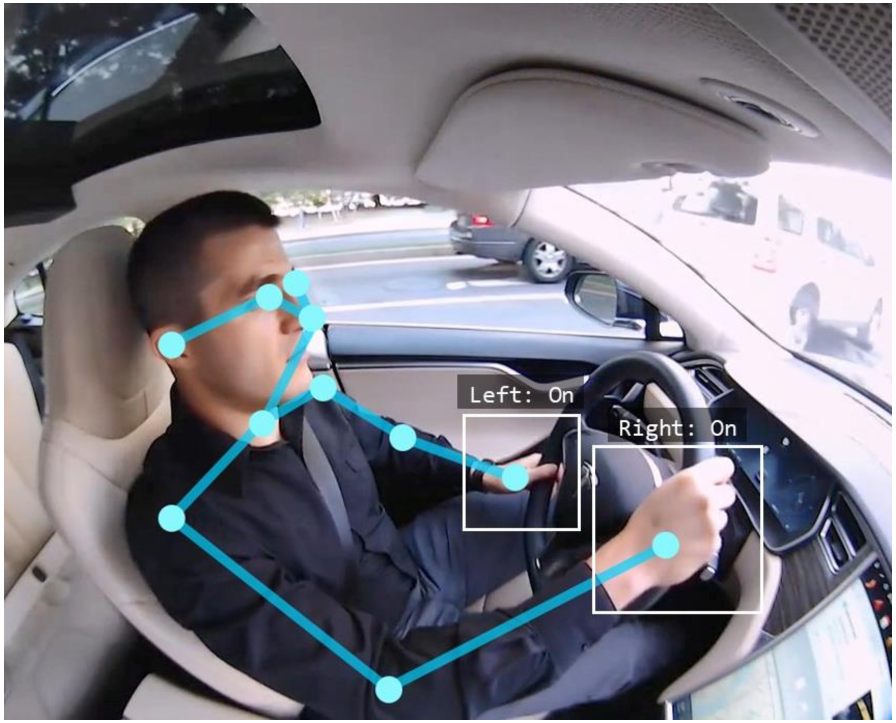

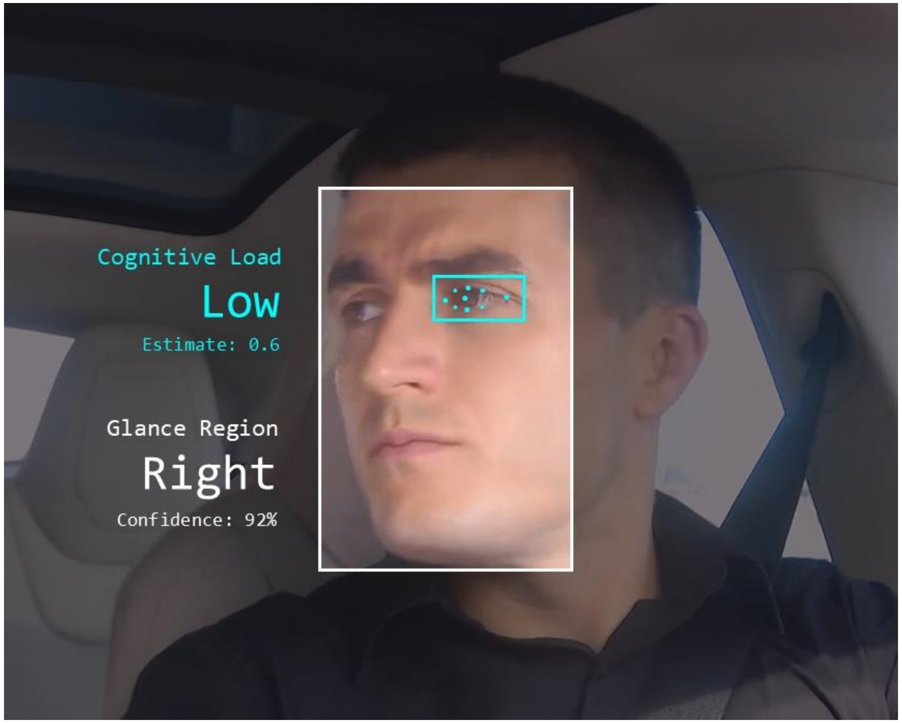

- Pattern of body movement
  - Vertical position in seat
  - General movement
- Beyond body movemnet
  - Smartphone
  - Hands on wheel
  - Activity
  - Context for DeepGlance

Sequential Detection Approach

Sequential Upper Body Pose Estimation:

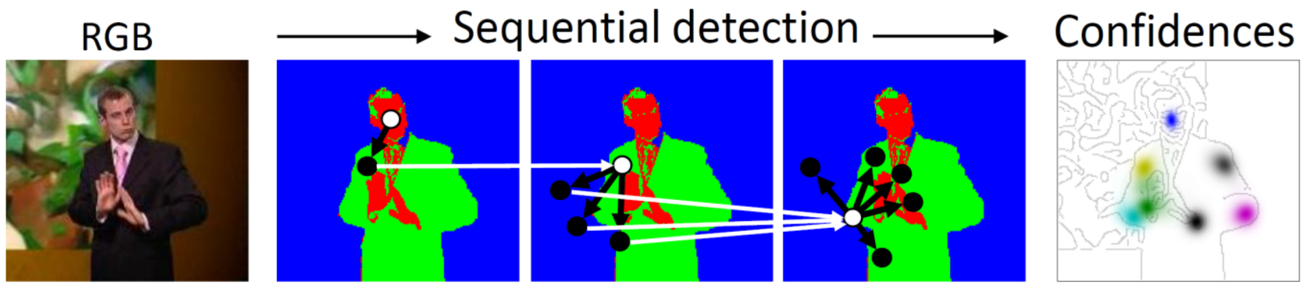

Temporal Fusion of Localized Confidences:

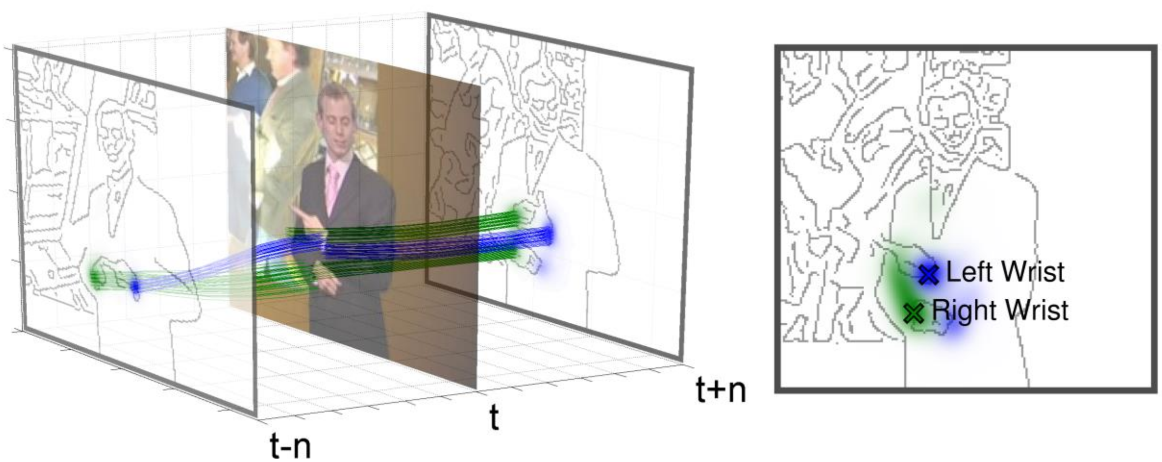

Charles, James, et al. "Upper body pose estimation with temporal sequential forests." *Proceedings of the British Machine Vision Conference 2014*. BMVA Press, 2014.

DeepPose: Holistic View

- Why holistic reasoning?
  - Besides extreme variability in articulations, **many of the joints are barely visible**

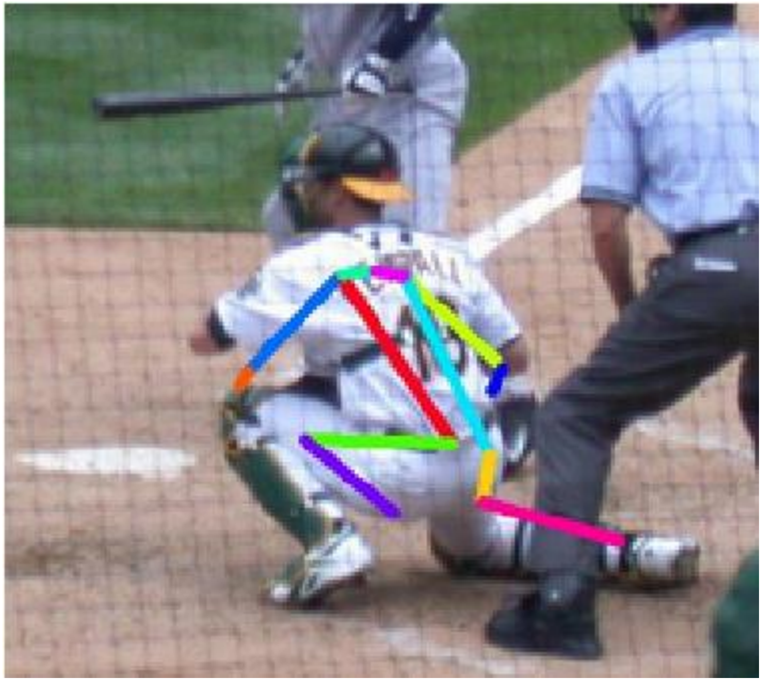

Cascade of Pose Regressors

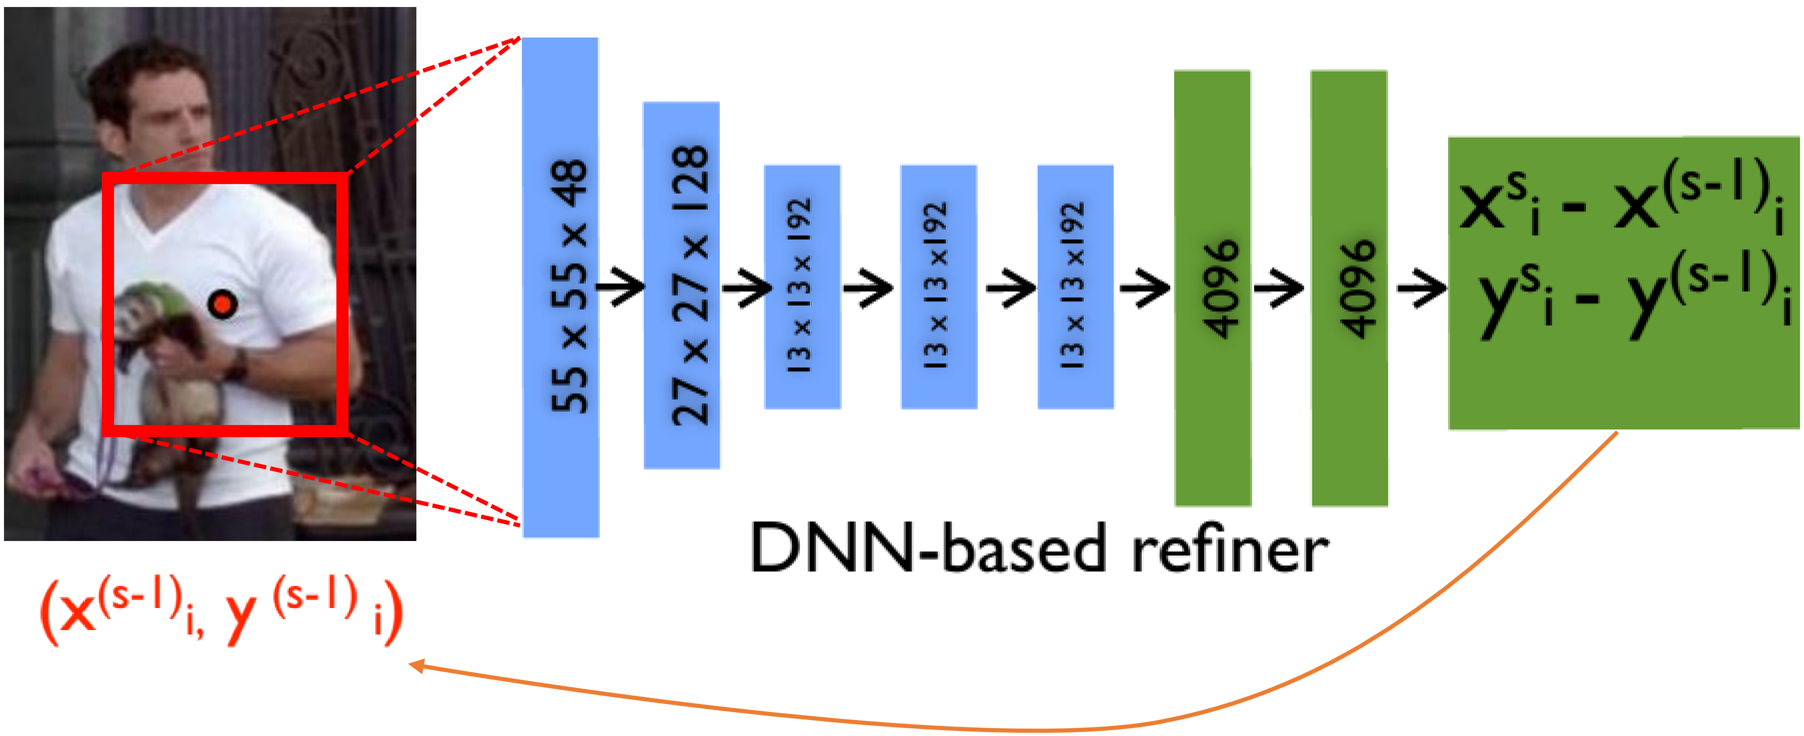

Part Detection

(a) Input image (b) Confidence maps

Assemble Parts: Part Affinity Fields

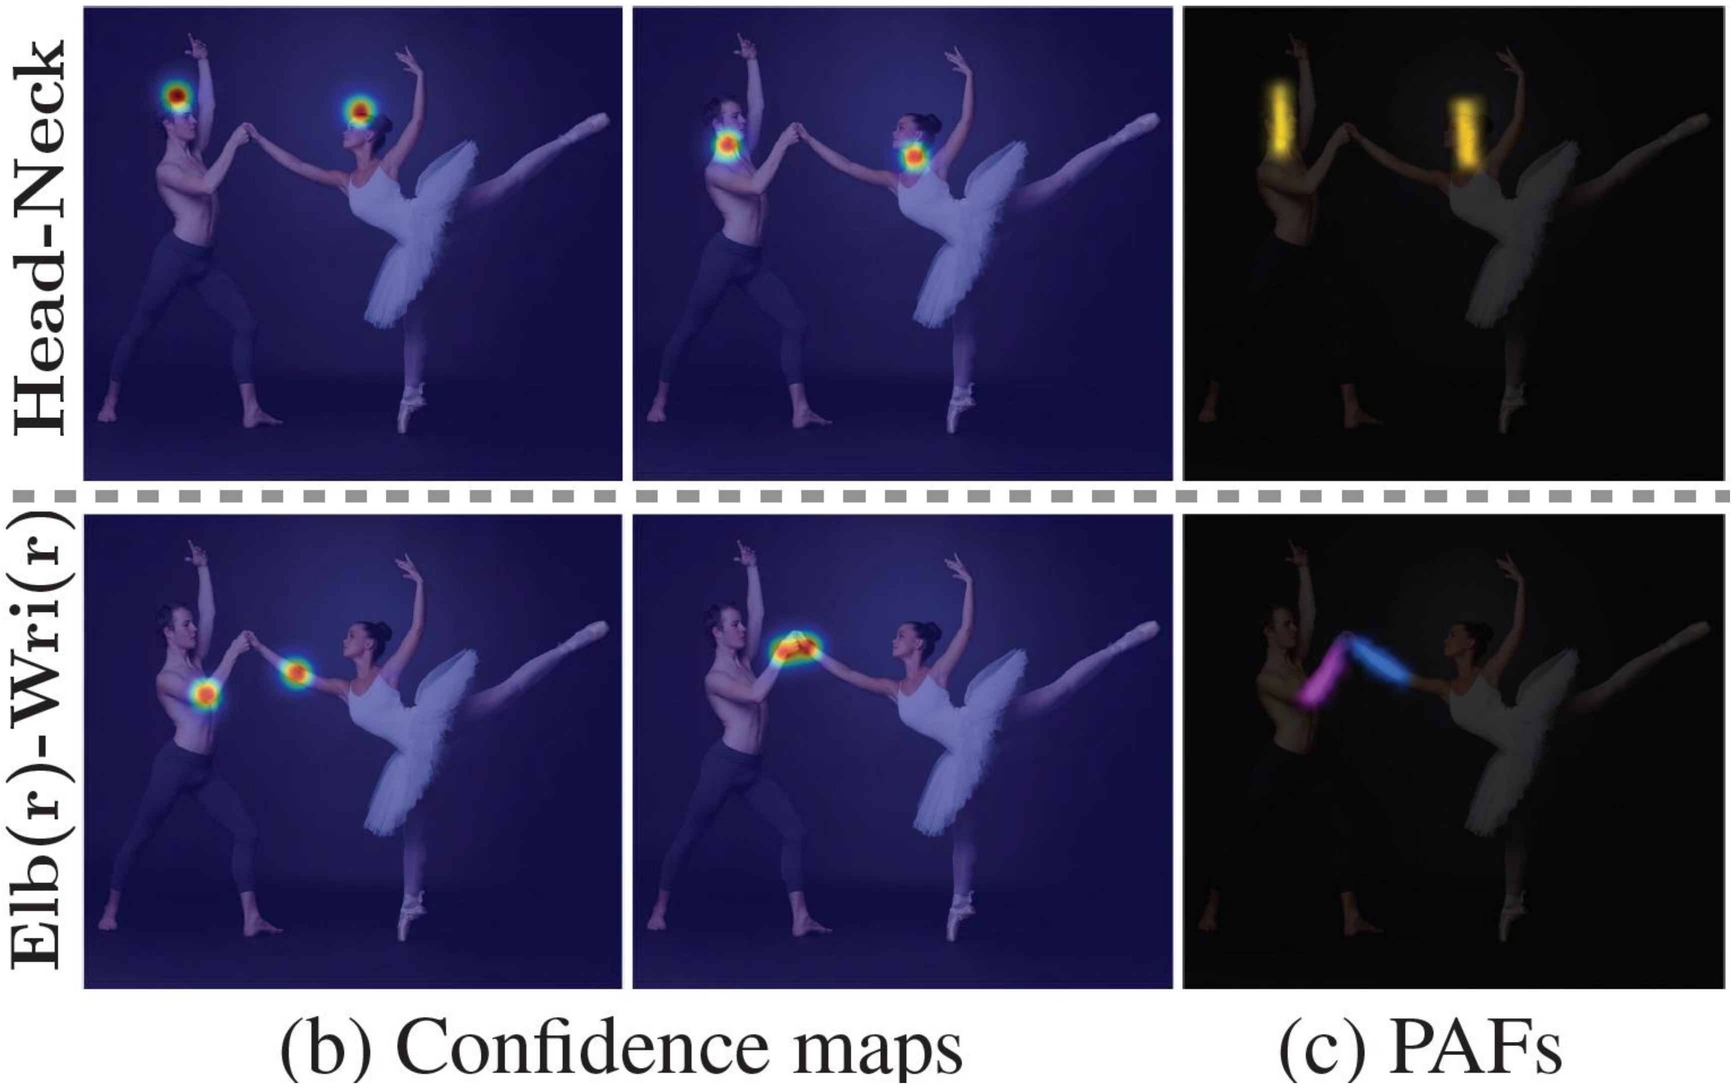

Bipartite Matching

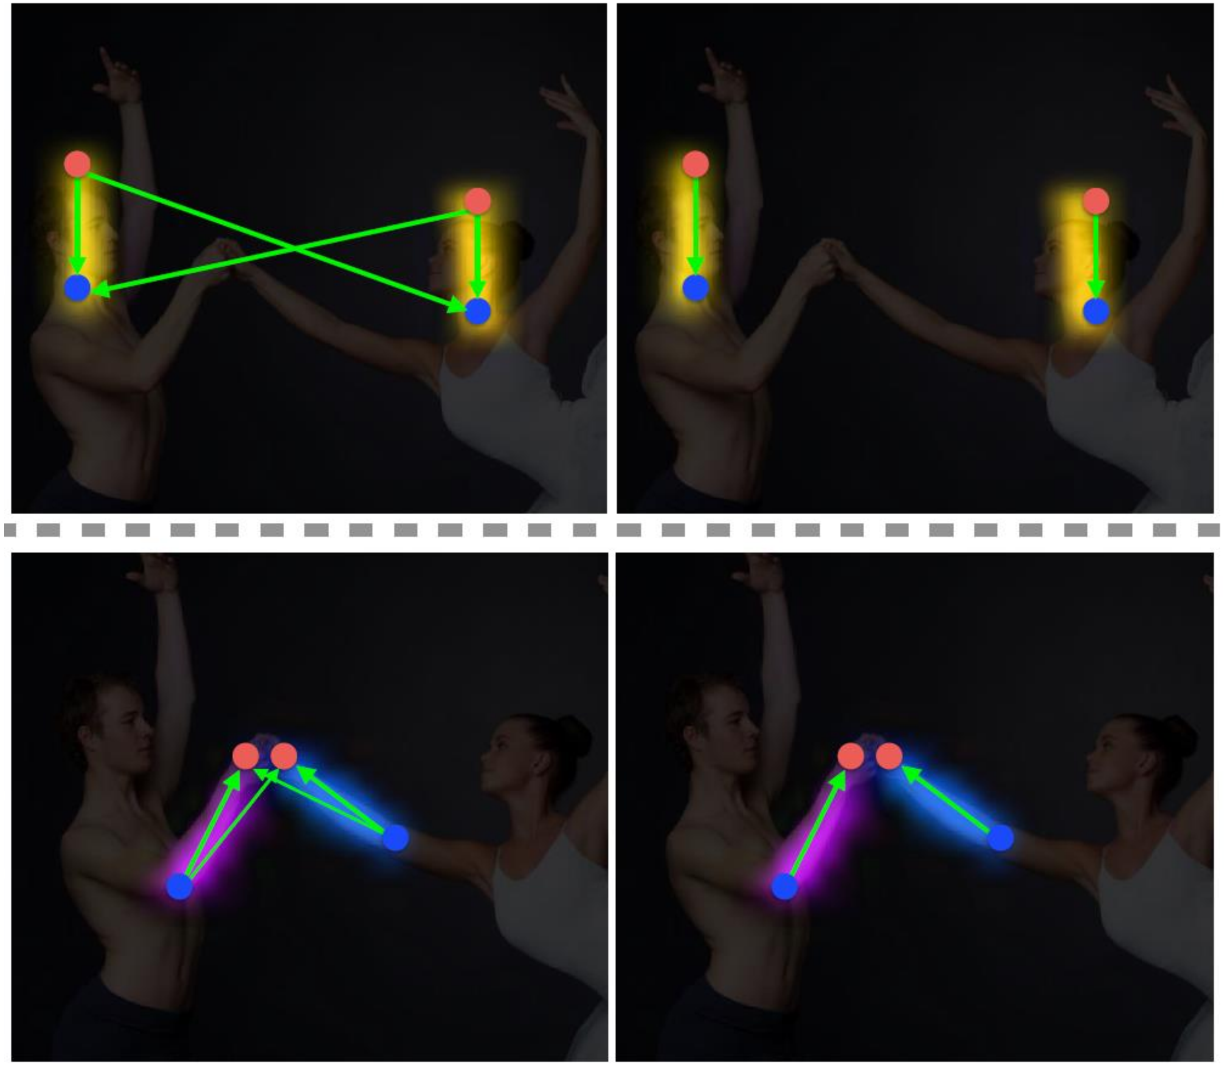

Temporal Convolutional Neural Networks

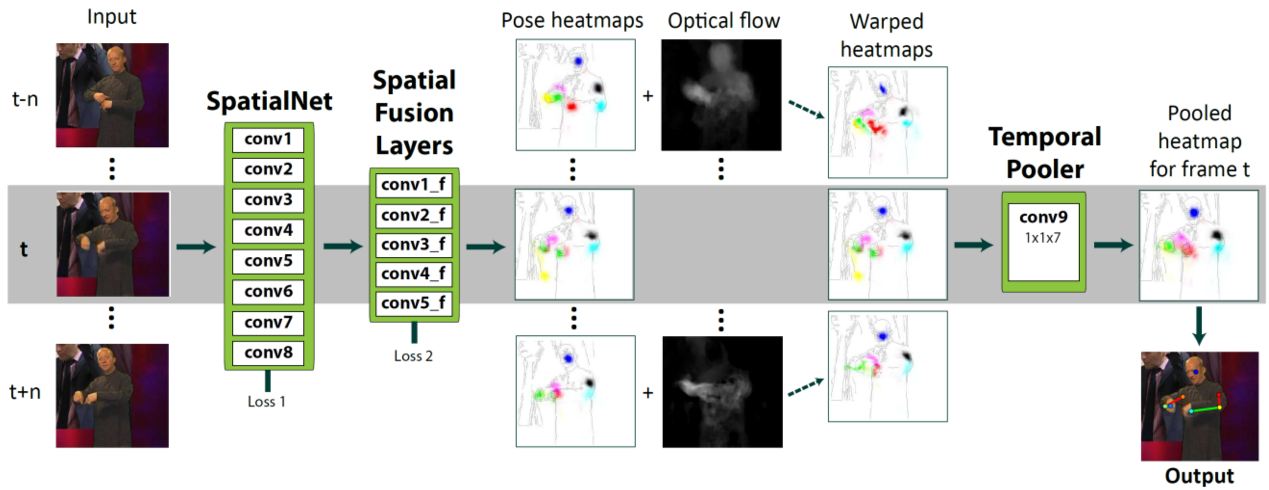

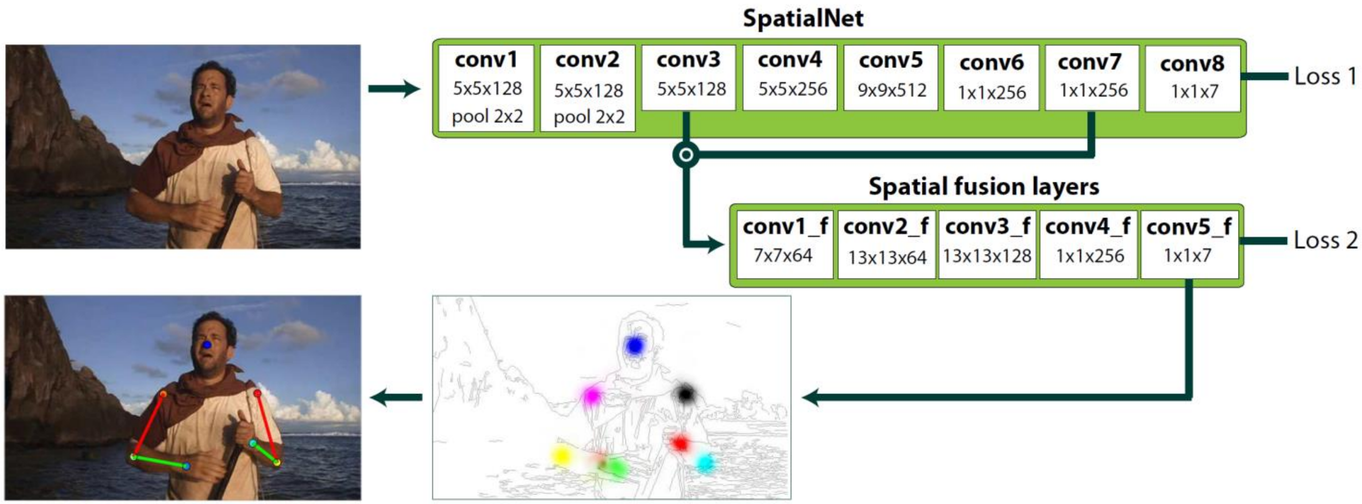

Pfister, Tomas, James Charles, and Andrew Zisserman. "Flowing convnets for human pose estimation in videos." *Proceedings of the IEEE International Conference on Computer Vision*. 2015.

Body Pose Estimation

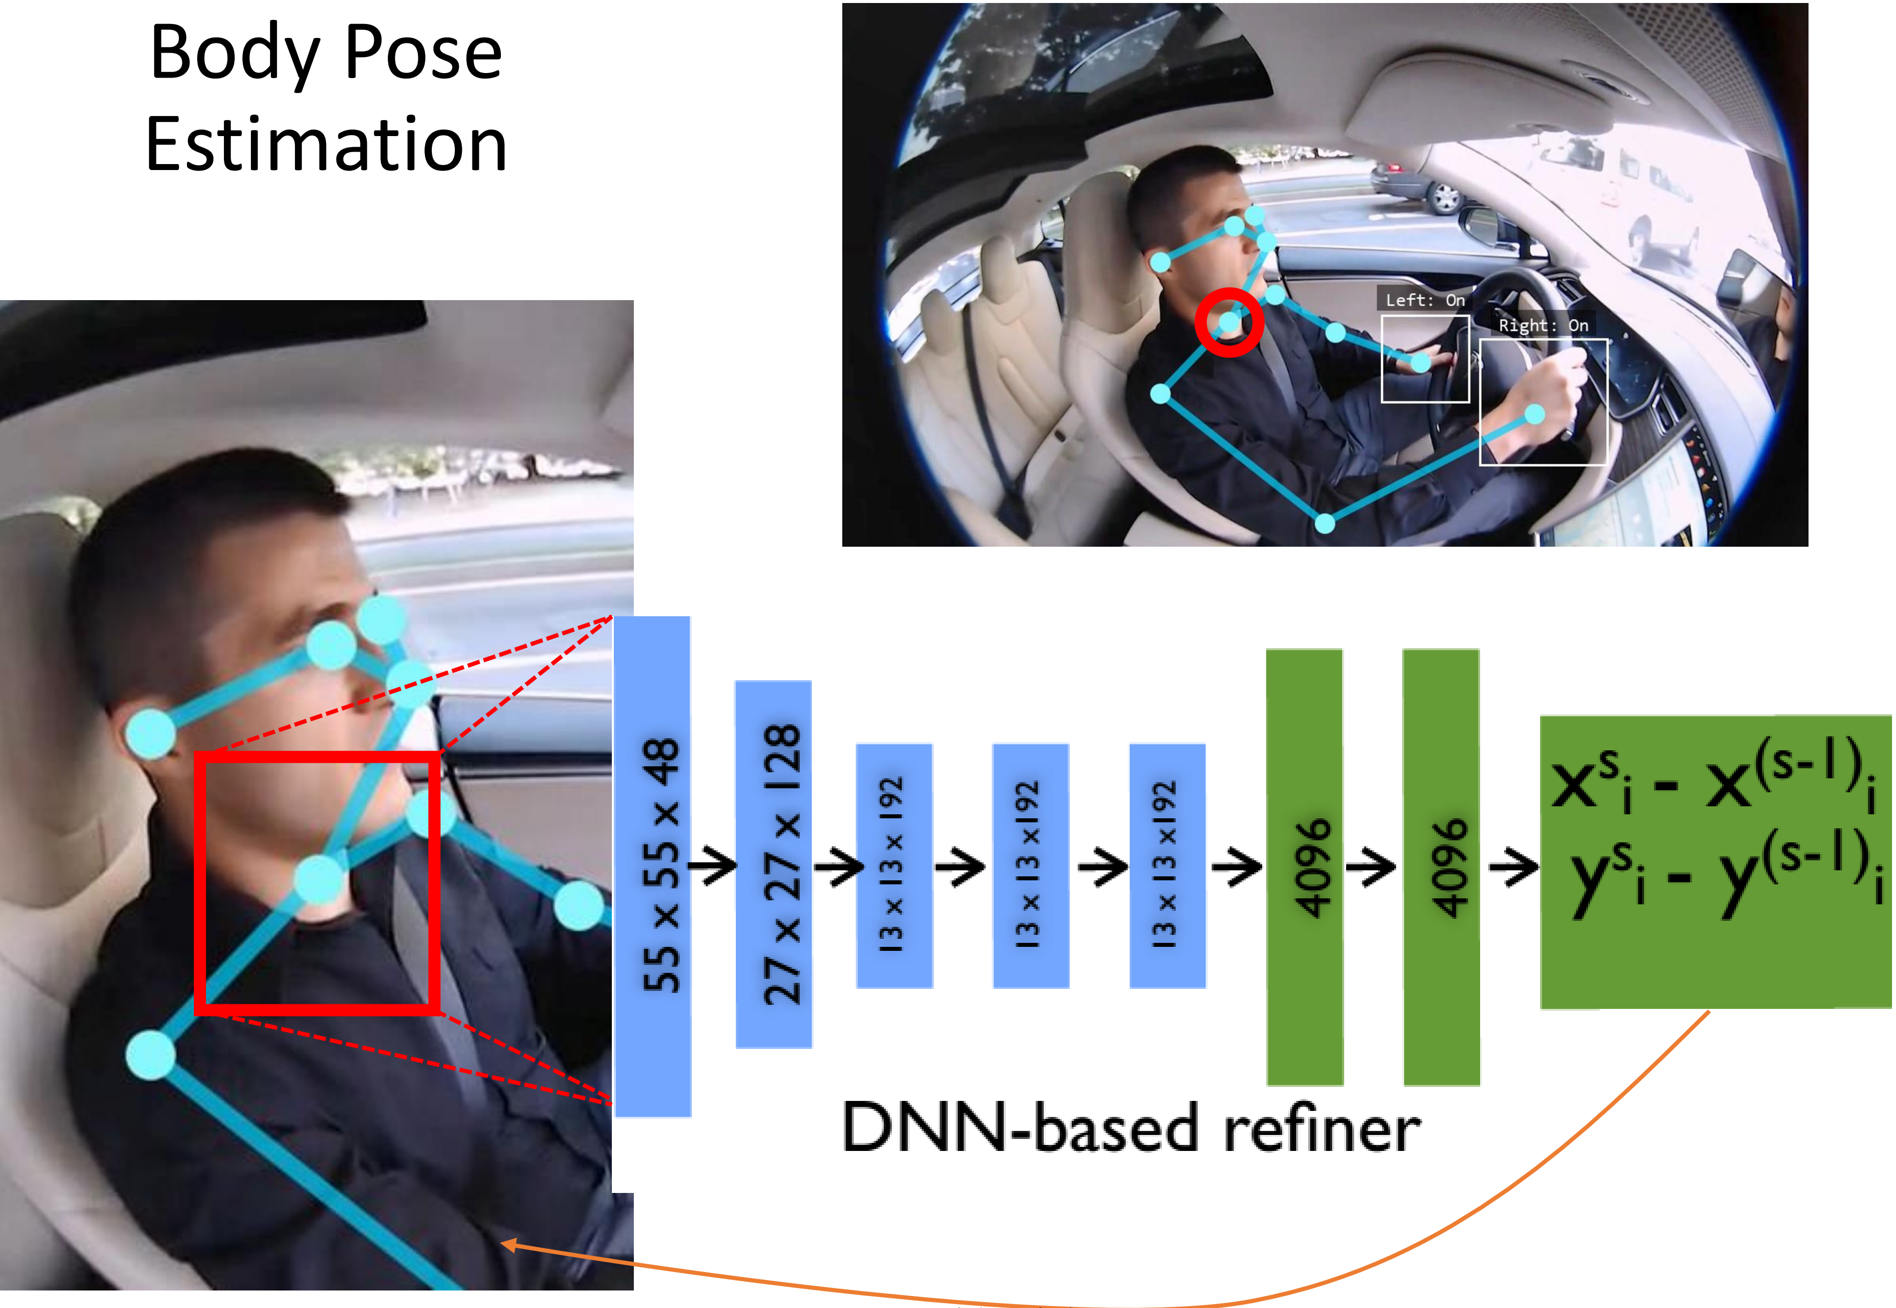

Body Pose: 20 Epochs (30 minutes each)

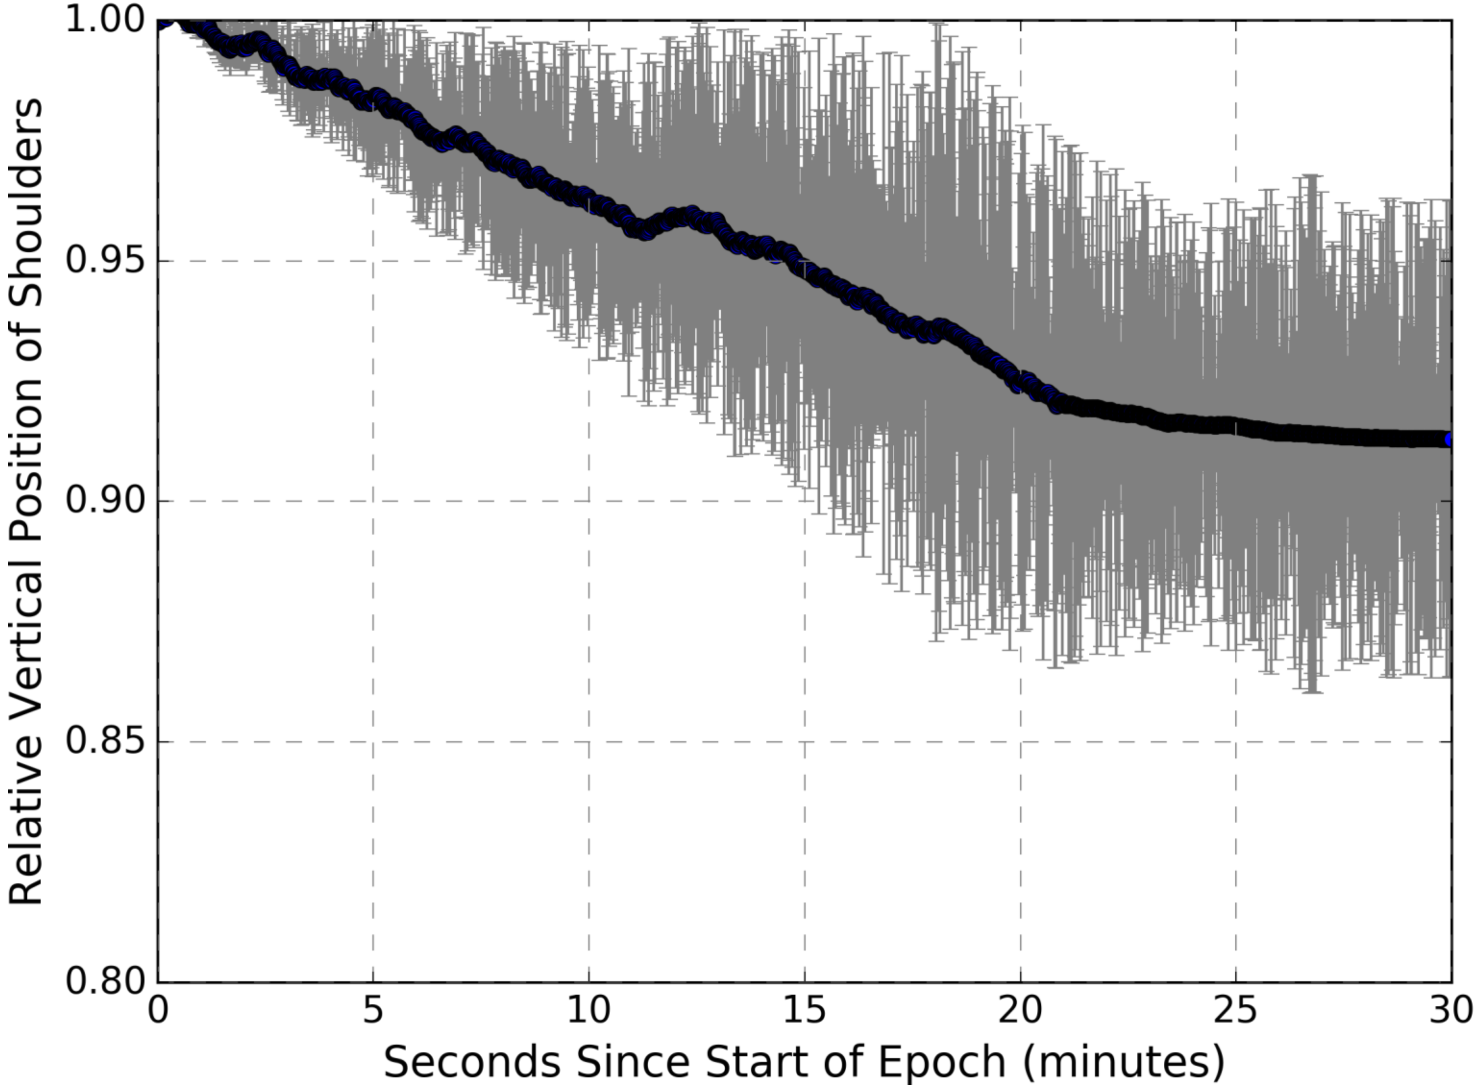

Body Pose: 20 Epochs (30 minutes each)

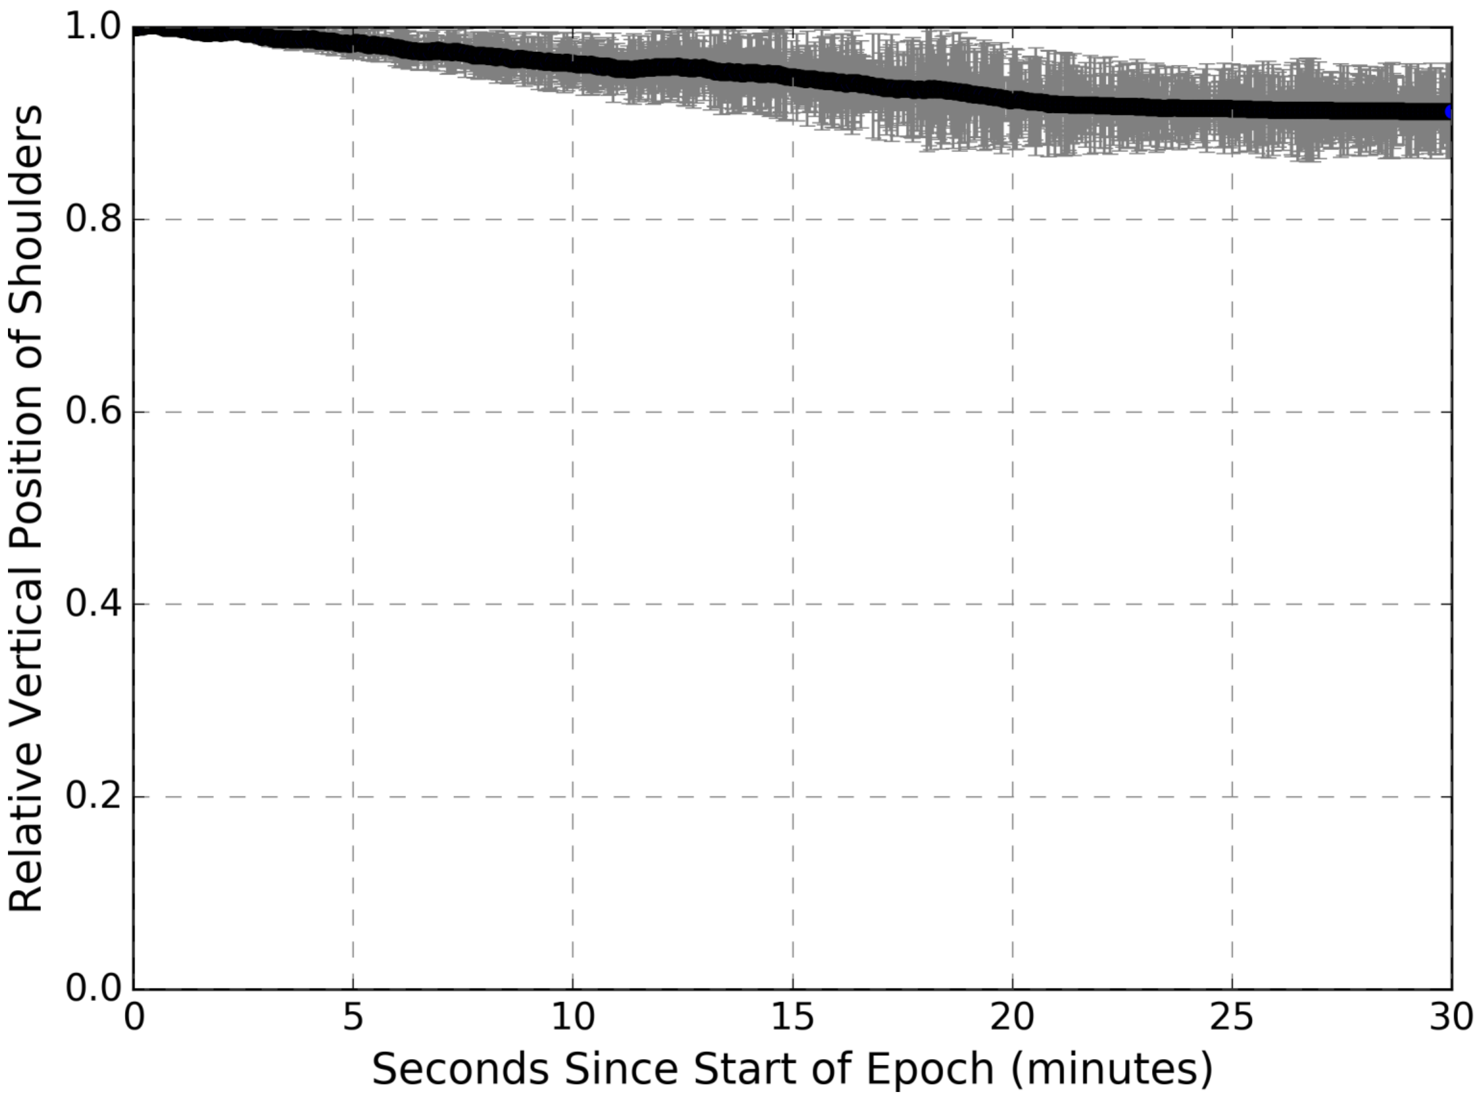

Pose Estimation
(Outside Vehicle Perspective)

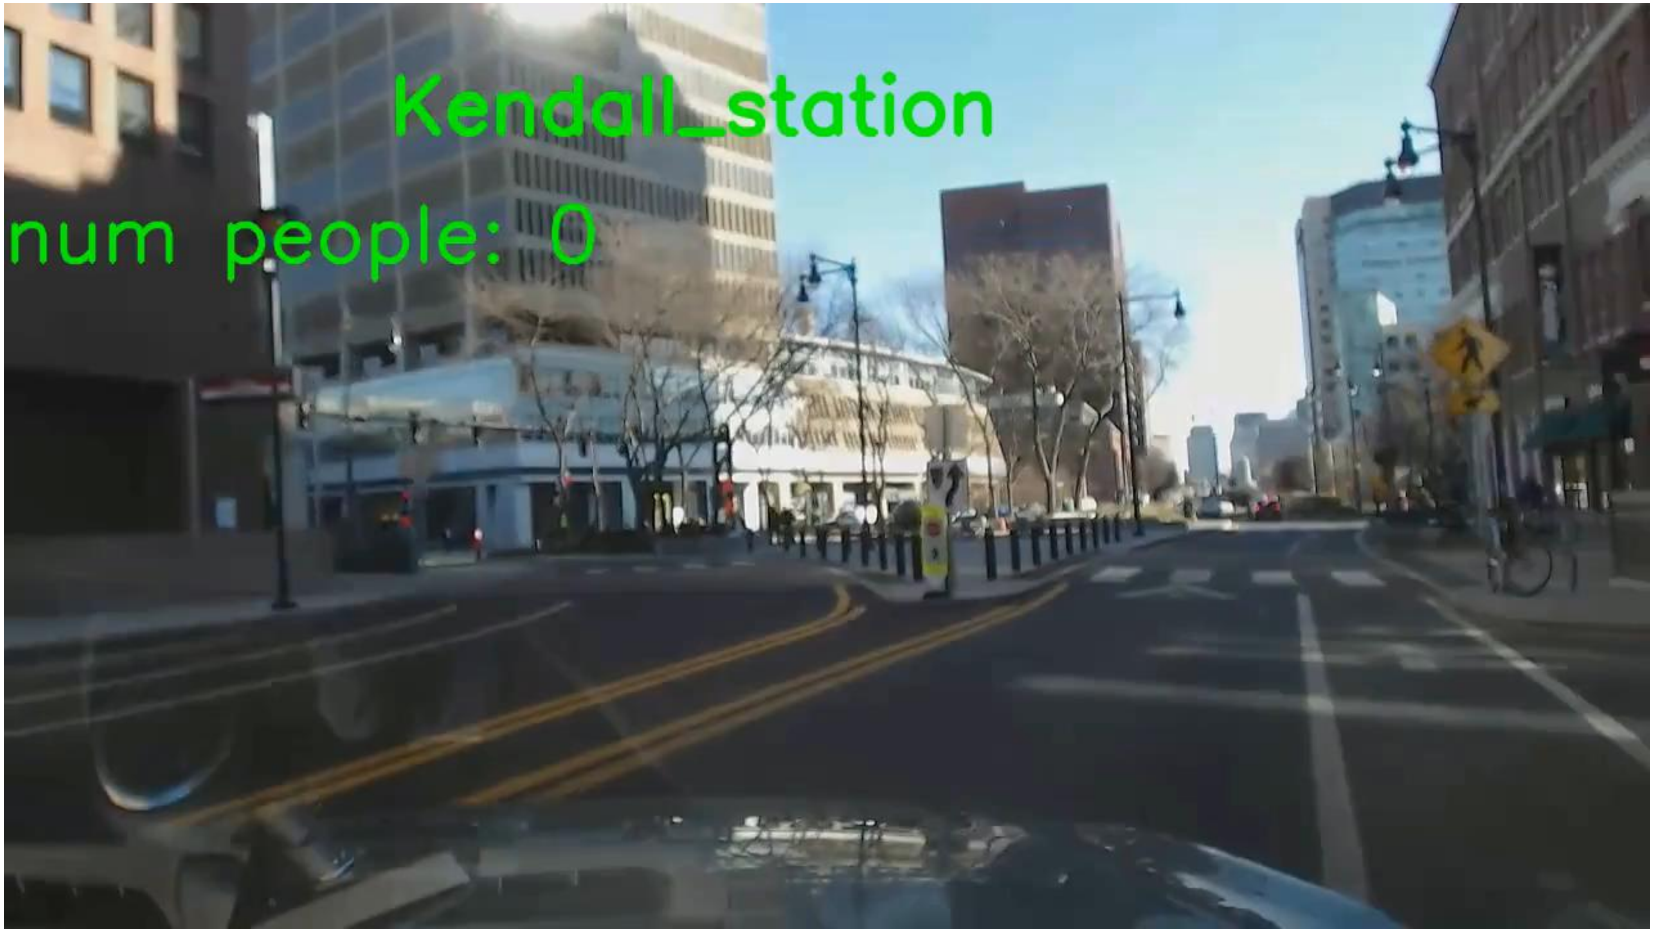

MIT Pedestrian Dataset

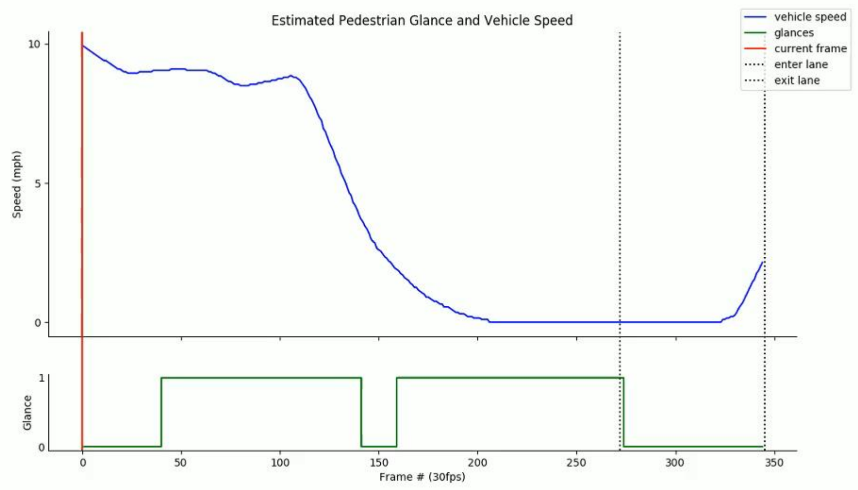

MIT Pedestrian Dataset

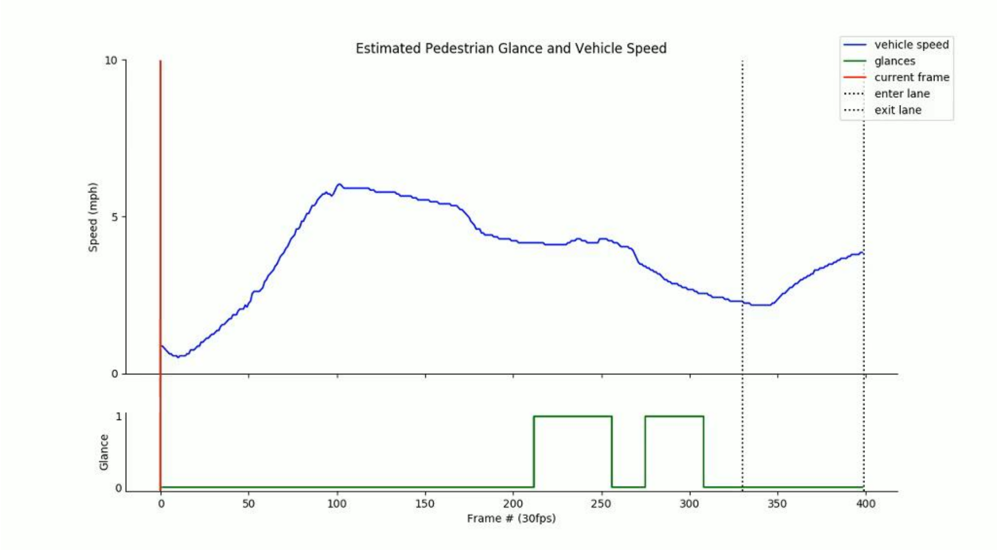

MIT Pedestrian Dataset

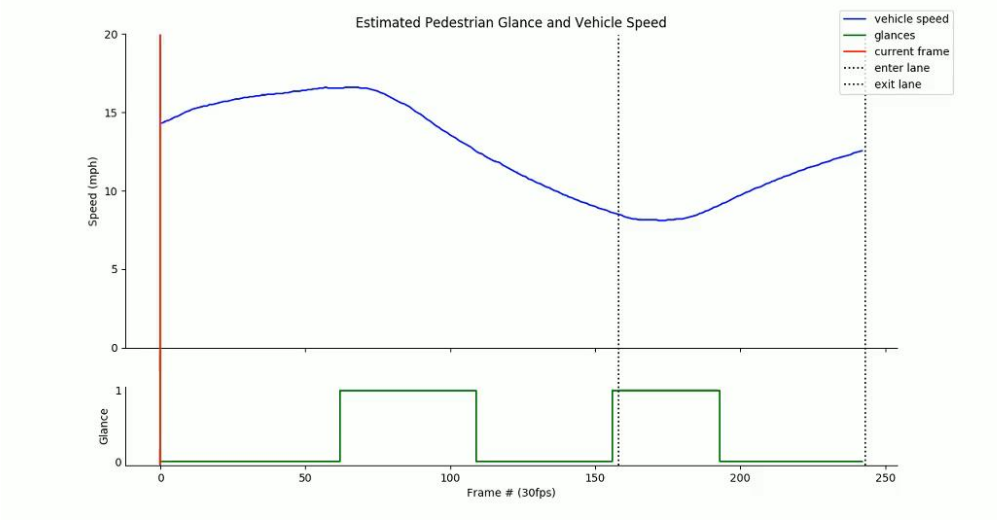

5rjs.cn 专注无人驾驶

Overview

- Human Imperfections
- Pedestrian Detection
- Body Pose Estimation
- **Glance Classification**
- Emotion Recognition
- Cognitive Load Estimation
- Human-Centered Vision for Autonomous Vehicles

Human Sensing:
A Deep Learning Perspective

Increasing level of detection resolution and **difficulty**

Pedestrian Detection | Body Pose | Head Pose | Blink Rate | Blink Duration | Eye Pose | Blink Dynamics | Pupil Diameter | Micro Saccades

Face Detection | Face Classification | Glance Classification | Drowsiness | Micro Glances | Cognitive Load

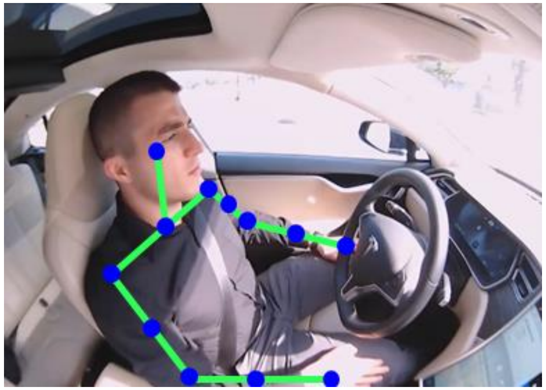

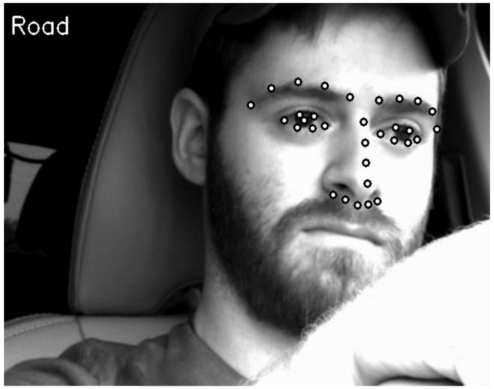

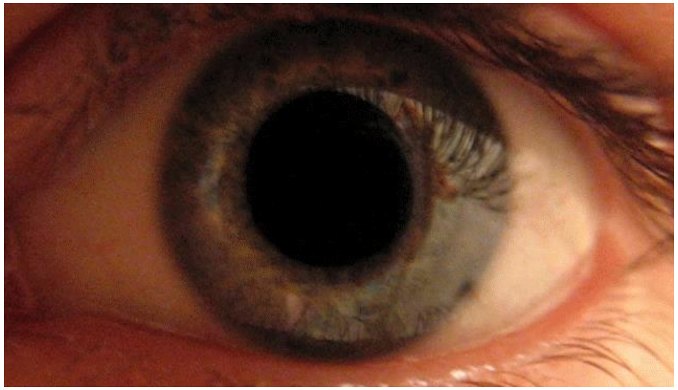

Glance Classification vs Gaze Estimation

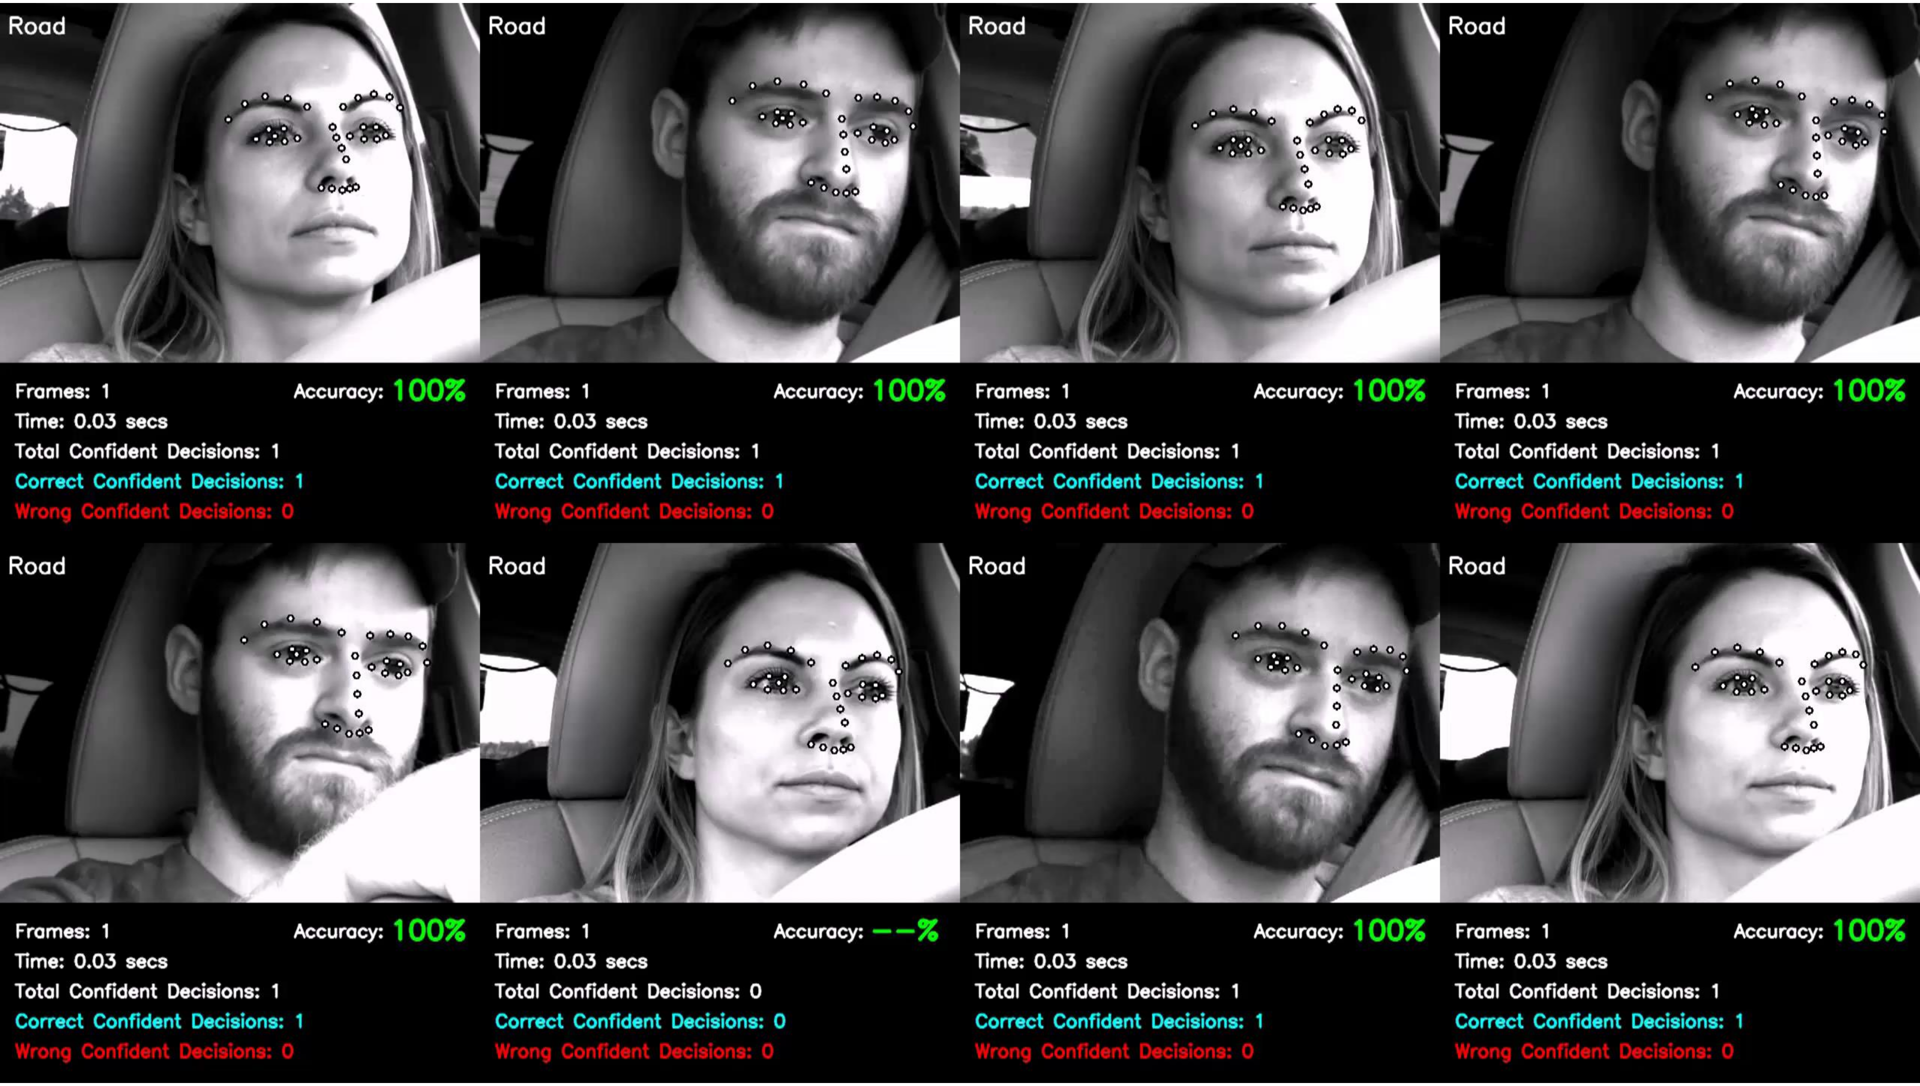

Pedestrian Glance Classification

Drive State Detection

- **Challenge:** real-world data is “messy”, have to deal with:
  - Vibration
  - Lighting variation
  - Body, head, eye movement

- **Solution:**
  - Automated calibration
  - Video stabilization (multi-resolutional)
  - Face part frontalization
  - Use deep neural networks (DNN)
    - No feature engineering
    - Use raw data

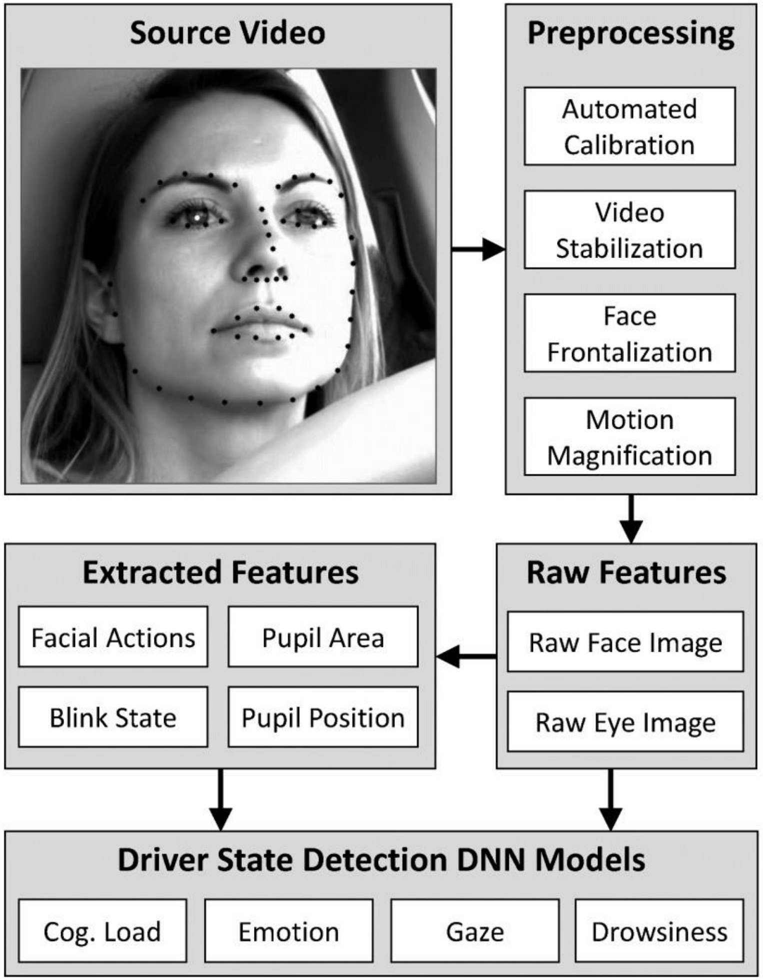

Face Alignment

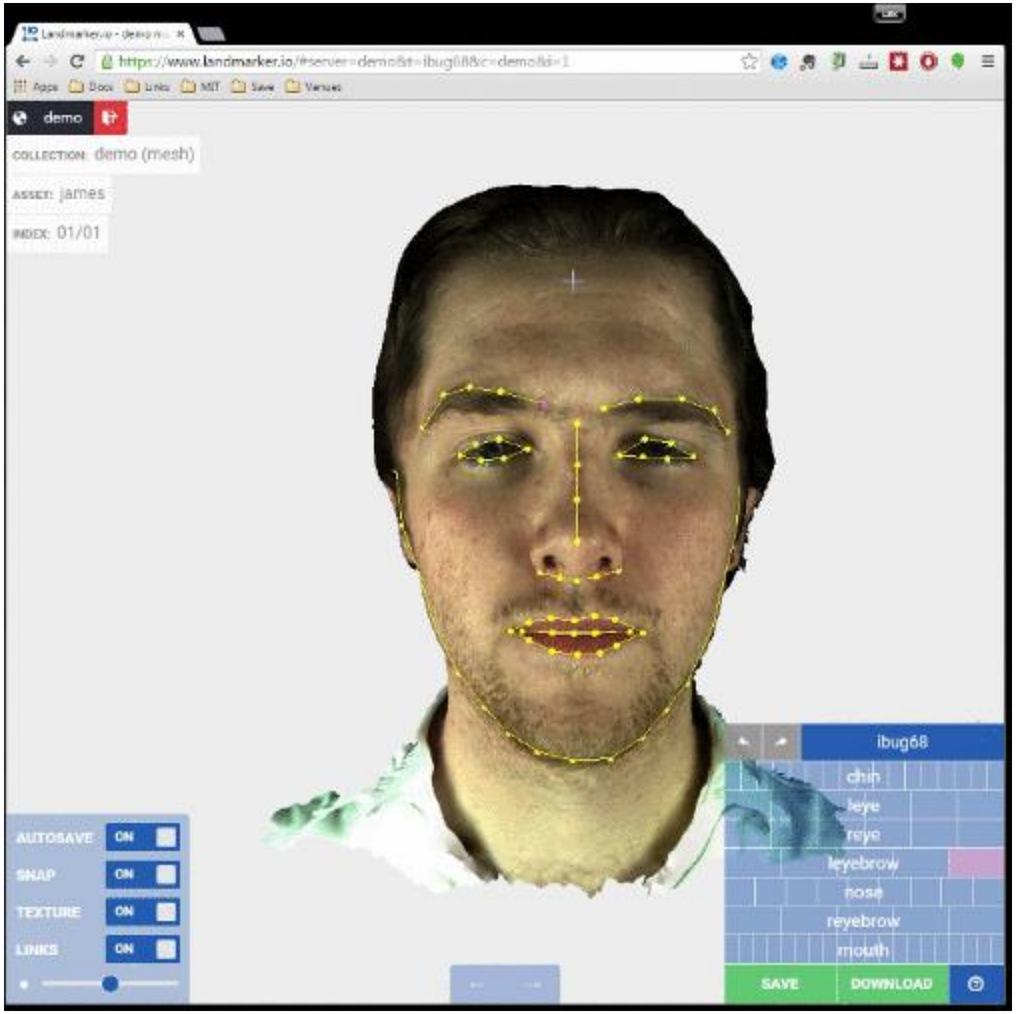

- Landmarker.io
  - Imperial College London
- Face in the Wild Challenge
  - XM2VTS
  - FRGC Ver.2
  - LFPW
  - HELEN
  - AFW
  - IBUG
- New Datasets
  - MPIIGaze
  - Columbia Gaze
  - 300VW

Gaze Classification Pipeline

1. Face detection *(the only easy step)*
2. Face alignment *(active appearance models or deep nets)*
3. Eye/pupil detection *(are the eyes visible?)*
4. Head (and eye) pose estimation *(+ normalization)*
5. Classification *(supervised learning = improves from data)*
6. Decision pruning *(how confident is the prediction)*

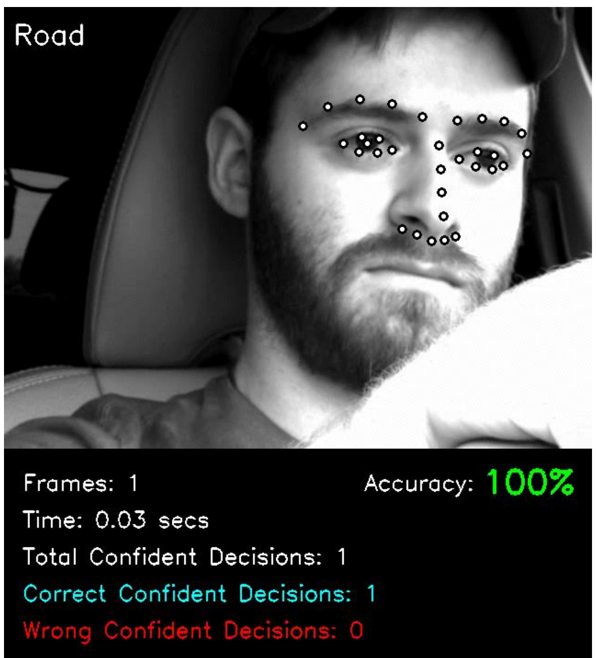

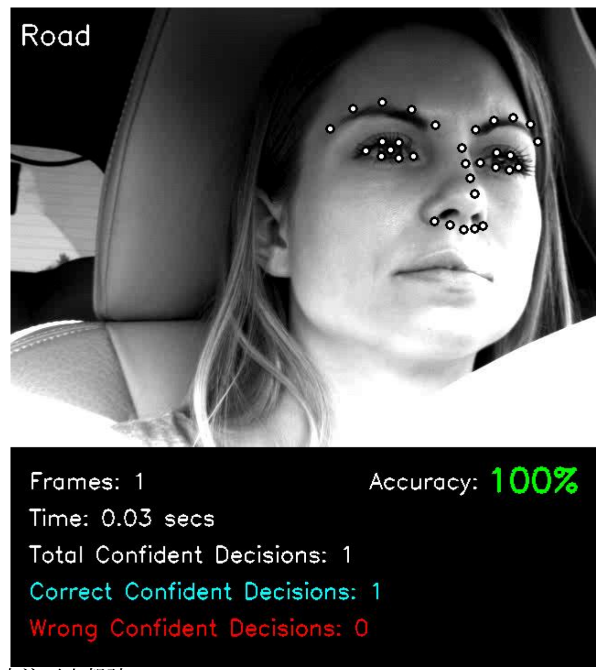

Annotation Tooling

"Semi-automated":

Ask a human for help with annotation when the machine is not confident.

Partial light occlusion

Full light occlusion

Move out of frame

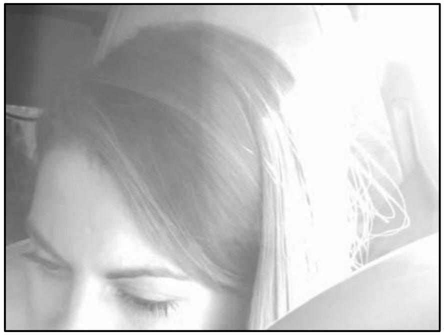

Hand occlusion

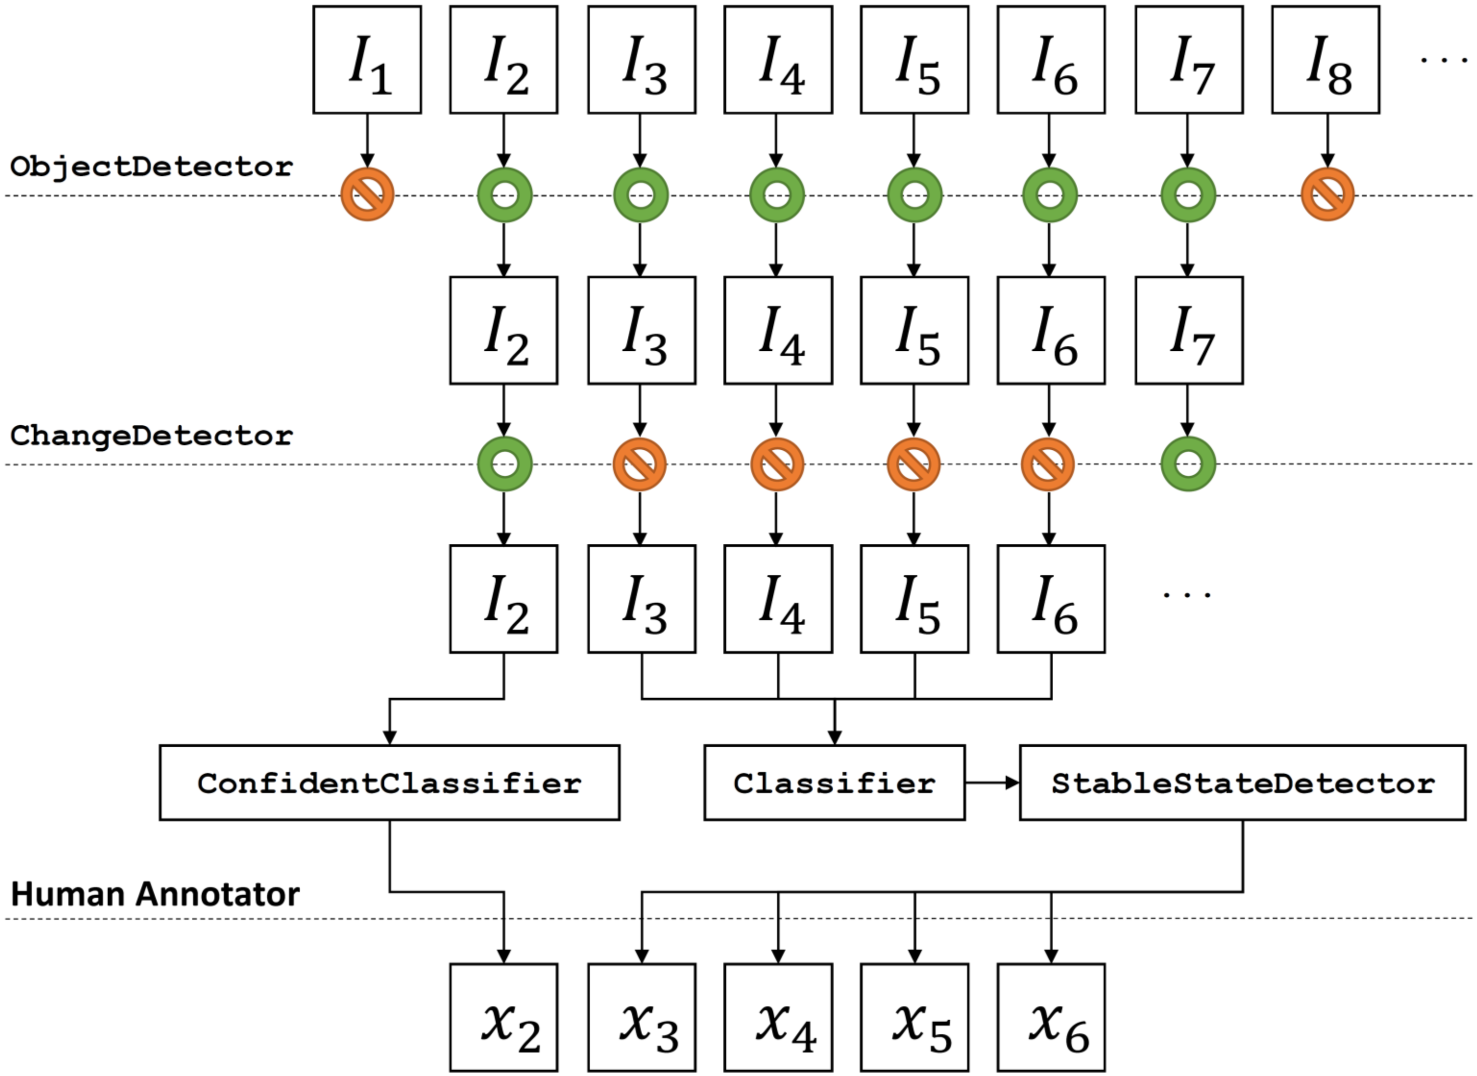

I_1 I_2 I_3 I_4 I_5 I_6 I_7 I_8 ...
ObjectDetector
I_2 I_3 I_4 I_5 I_6 I_7
ChangeDetector
I_2 I_3 I_4 I_5 I_6 ...
ConfidentClassifier
Classifier
StableStateDetector
Human Annotator
x_2 x_3 x_4 x_5 x_6

Semi-Automated Annotation Work Flow

* Human in **red** and machine in **blue**

1. Select and load in video of driver face.
2. Detect face: have we seen this person before?
3. Localize camera: have we seen this angle before?
4. Provide tradeoff between accuracy and percent frames.
5. Select target accuracy: 95%, 99%, or 99.9%
6. Perform gaze classification on full video *(1 hour per 1 hour of video)*
7. Step through and annotate the frames machine did not classify.
8. (Optional) Re-run steps 6 and 7.
9. Enjoy fully annotated video!

Real-Time Glance Classification

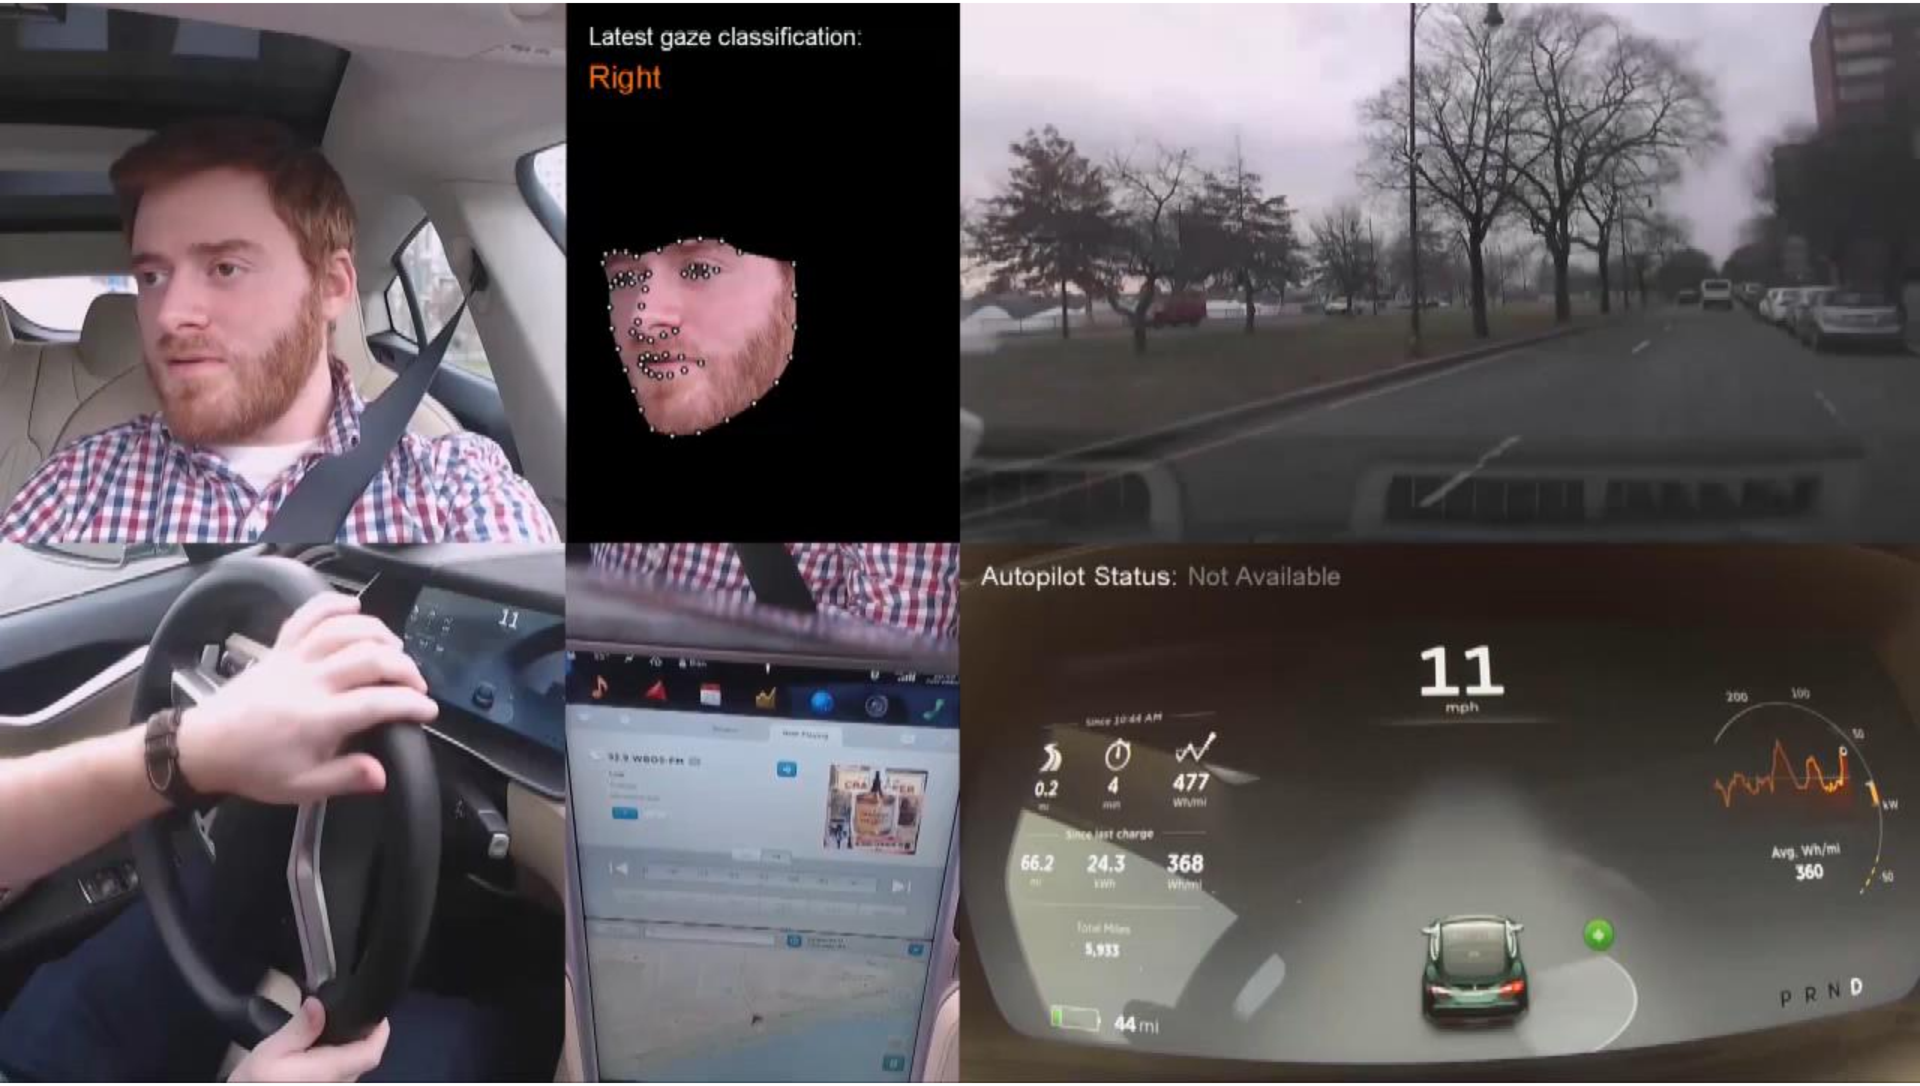

Overview

- Human Imperfections
- Pedestrian Detection
- Body Pose Estimation
- Glance Classification
- **Emotion Recognition**
- Cognitive Load Estimation
- Human-Centered Vision for Autonomous Vehicles

Human Sensing:
A Deep Learning Perspective

Increasing level of detection resolution and **difficulty**

Pedestrian Detection | Body Pose | Head Pose | Blink Rate | Blink Duration | Eye Pose | Blink Dynamics | Pupil Diameter | Micro Saccades

Face Detection | Face Classification | Glance Classification | Drowsiness | Micro Glances | Cognitive Load

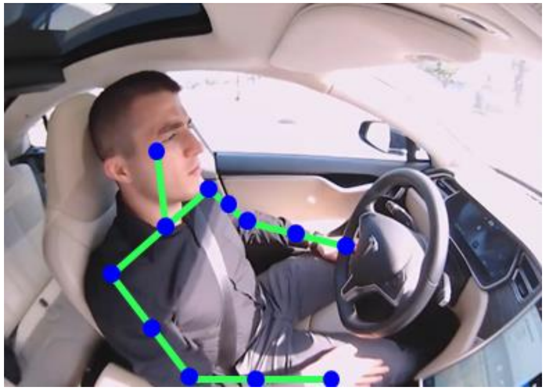

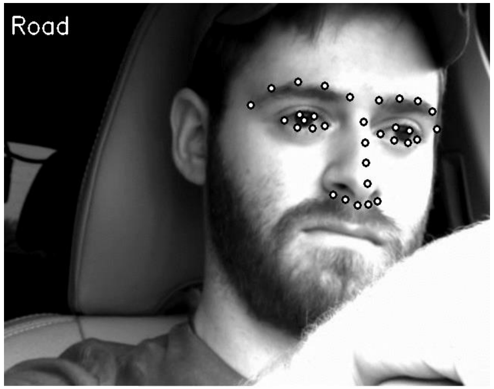

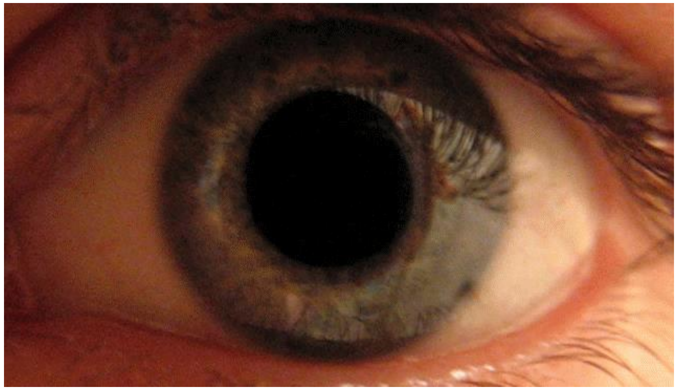

Drive State Detection

- **Challenge:** real-world data is “messy”, have to deal with:
  - Vibration
  - Lighting variation
  - Body, head, eye movement
- **Solution:**
  - Automated calibration
  - Video stabilization (multi-resolutional)
  - Face part frontalization
  - Use deep neural networks (DNN)
    - No feature engineering
    - Use raw data

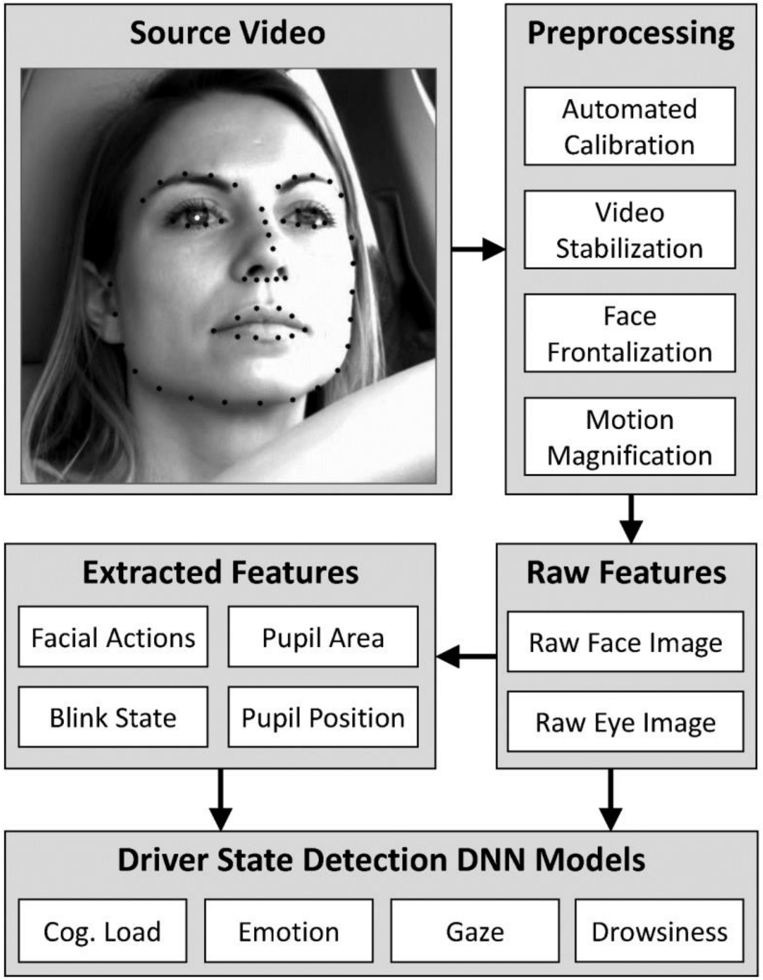

Emotion Recognition

- Many ways to taxonomize emotion.
- Example:
  Parrot's primary emotions:
  - Love
  - Joy
  - Surprise
  - Anger
  - Sadness
  - Fear
- Two approaches
  - General
  - Application-specific

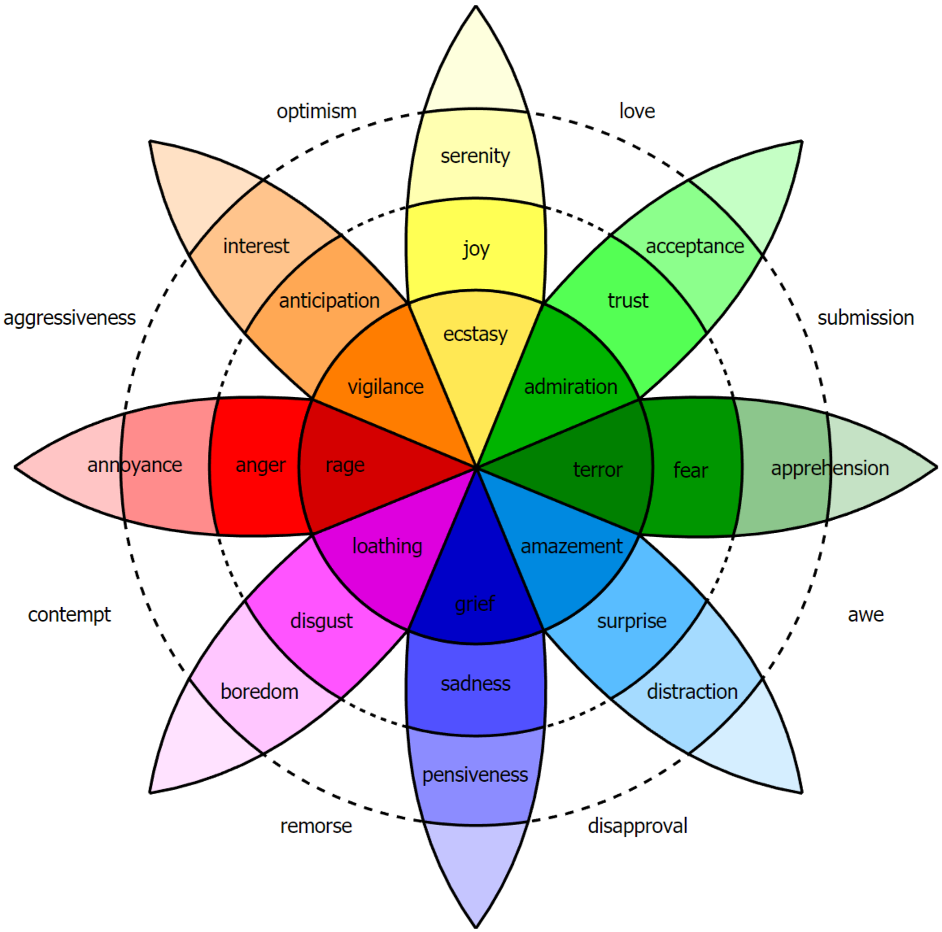

Building Blocks: Facial Expressions

- 42 individual facial muscles in the face.

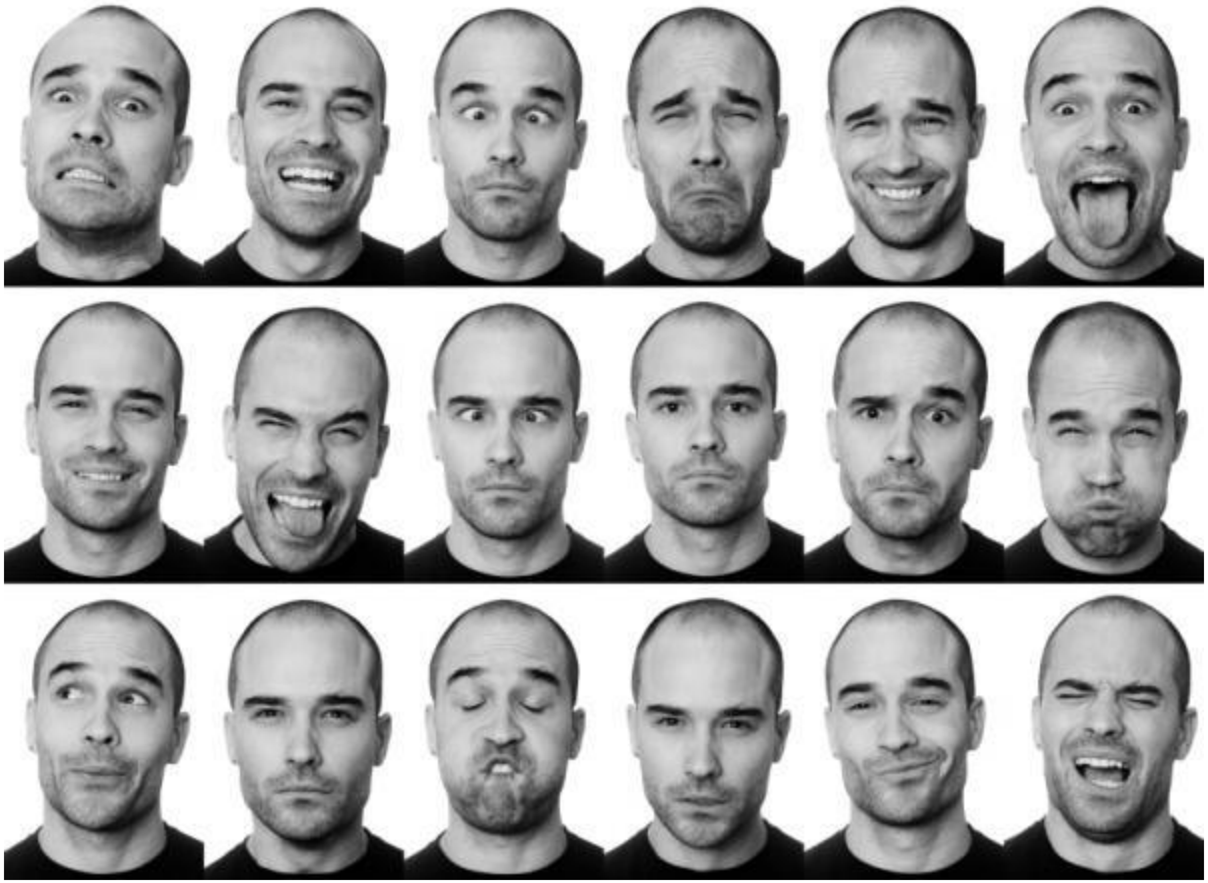

General Emotion Recognition

Example: Affectiva SDK

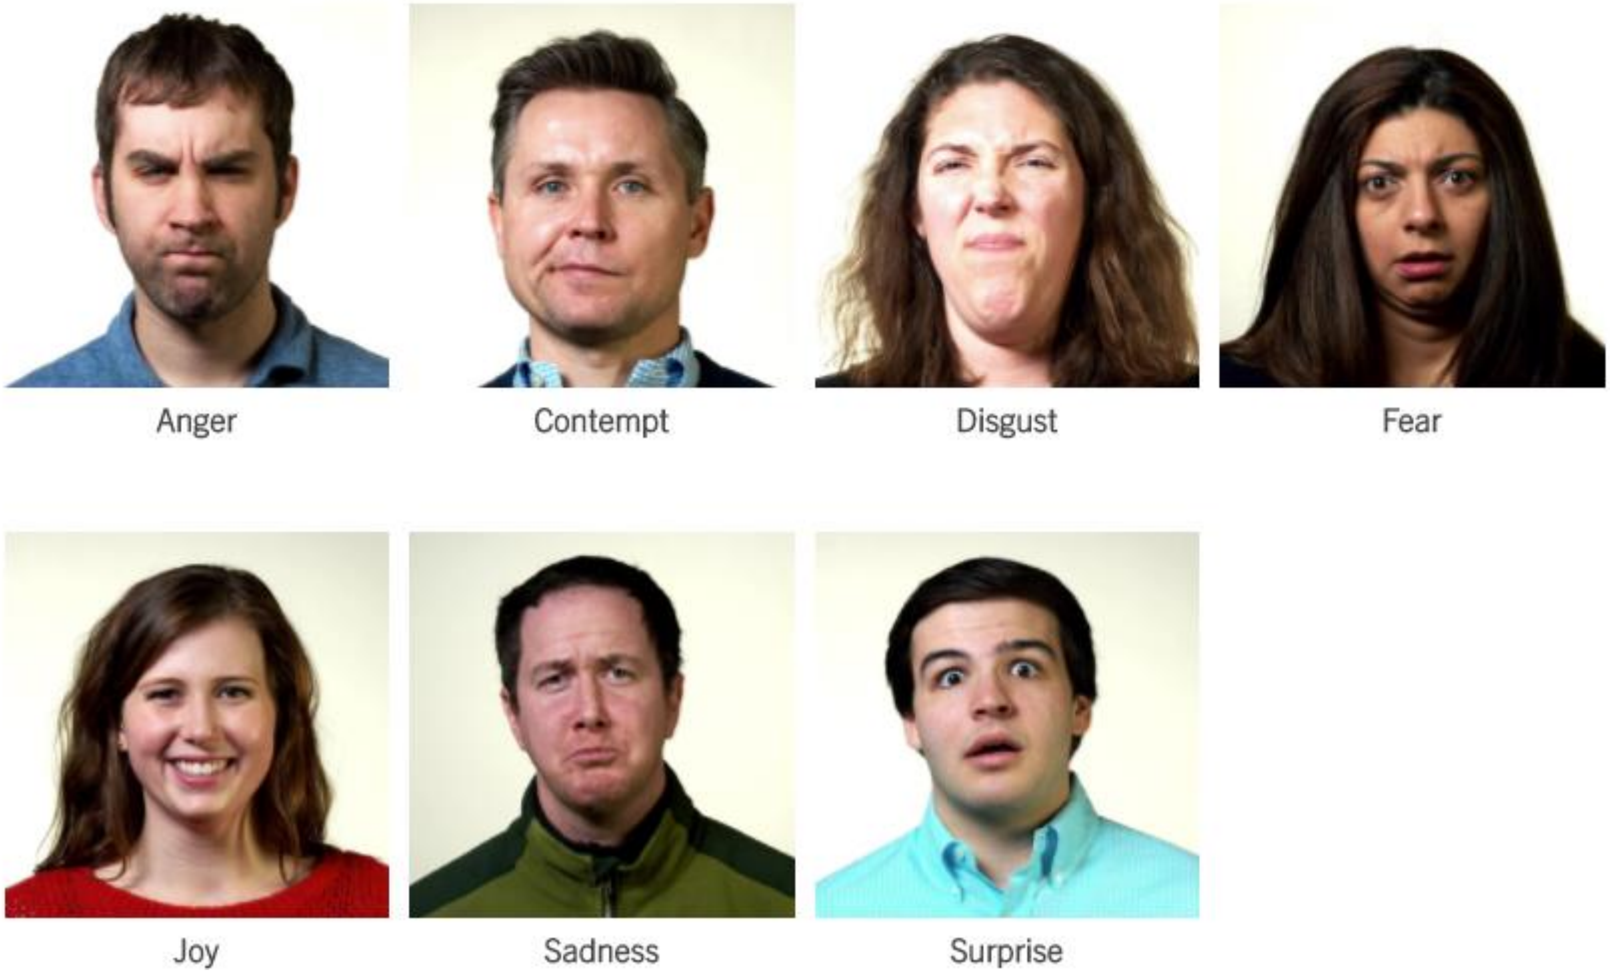

General Emotion Recognition

Example: Affectiva SDK

Emotion	Increase Likelihood	Decrease Likelihood
Joy	Smile	Brow Raise Brow Furrow
Anger	Brow furrow Lid Tighten Eye Widen Chin Raise Mouth Open Lip Suck	Inner Brow Raise Brow Raise Smile
Disgust	Nose Wrinkle Upper Lip Raise	Lip Suck Smile

Application-Specific Emotion Recognition: Driver **Frustration**

Class 1: Satisfied with Voice-Based Interaction

Class 2: Frustrated with Voice-Based Interaction

Emotion Generation

https://agi.mit.edu

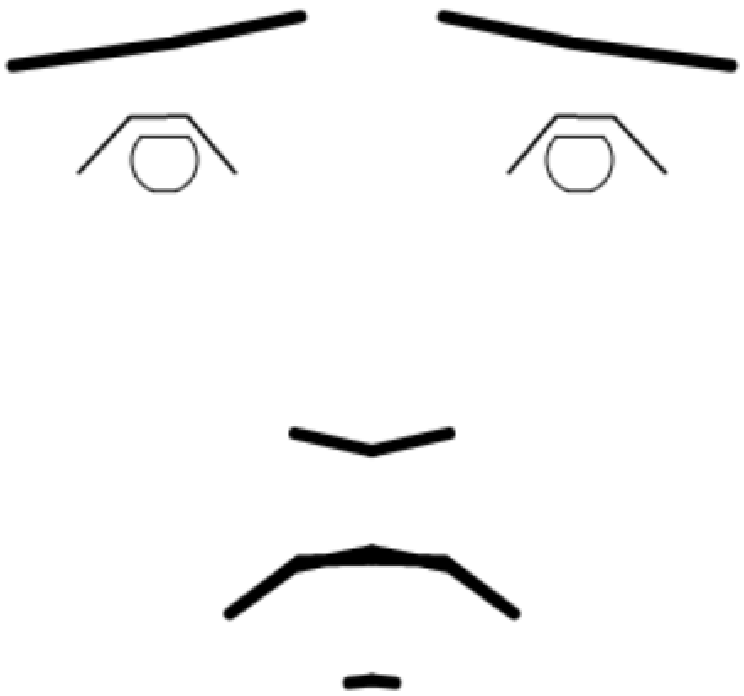

Overview

- Human Imperfections
- Pedestrian Detection
- Body Pose Estimation
- Face Detection
- Glance Classification
- Emotion Recognition
- **Cognitive Load Estimation**
- Human-Centered Vision for Autonomous Vehicles

Human Sensing:
A Deep Learning Perspective

Increasing level of detection resolution and **difficulty**

- Pedestrian Detection
- Body Pose
- Head Pose
- Blink Rate
- Blink Duration
- Eye Pose
- Blink Dynamics
- Pupil Diameter
- Micro Saccades
- Face Detection
- Face Classification
- Glance Classification
- Drowsiness
- Micro Glances
- Cognitive Load

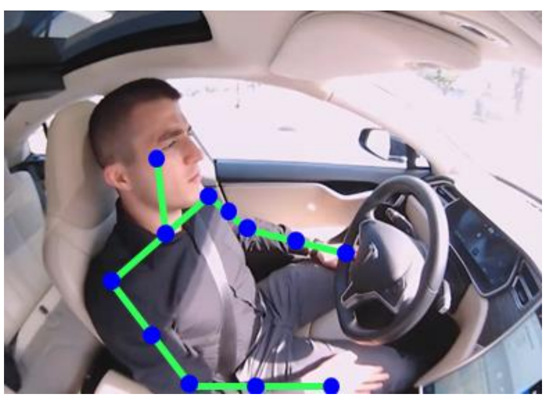

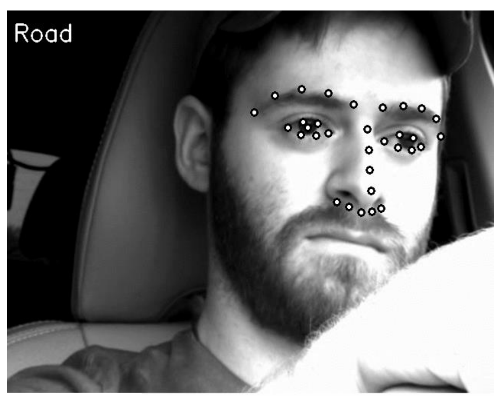

Eye in Motion: Saccades

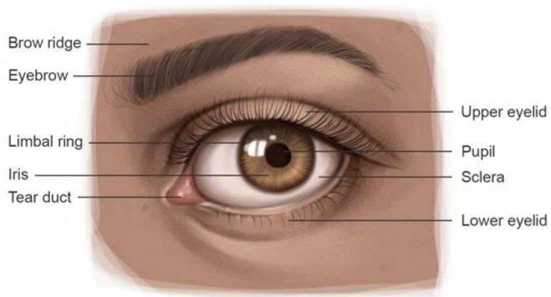

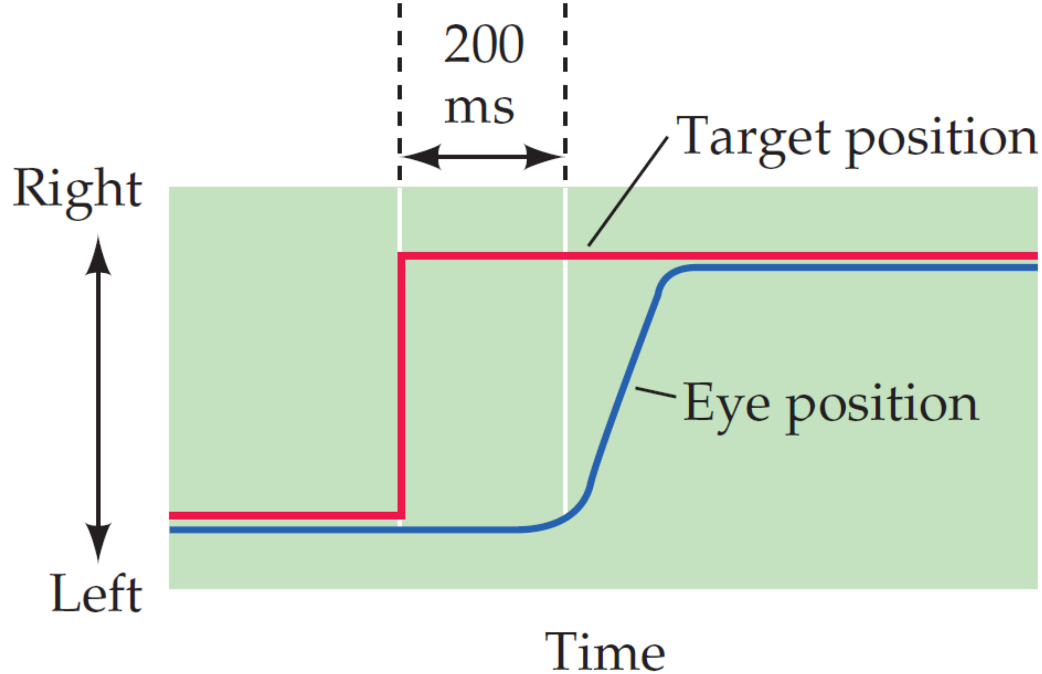

- Ballistic movements
- Can be small or large (reading vs exploring the room)
- Can be voluntary or reflexive
- During 200ms period: compute the position of target with respect to fovea and convert to motor command
- The eye movement is 15-100 ms
- If target moves during eye movement, adjustments have to be made **after** movement is completed.

Eye in Motion: Smooth Pursuits

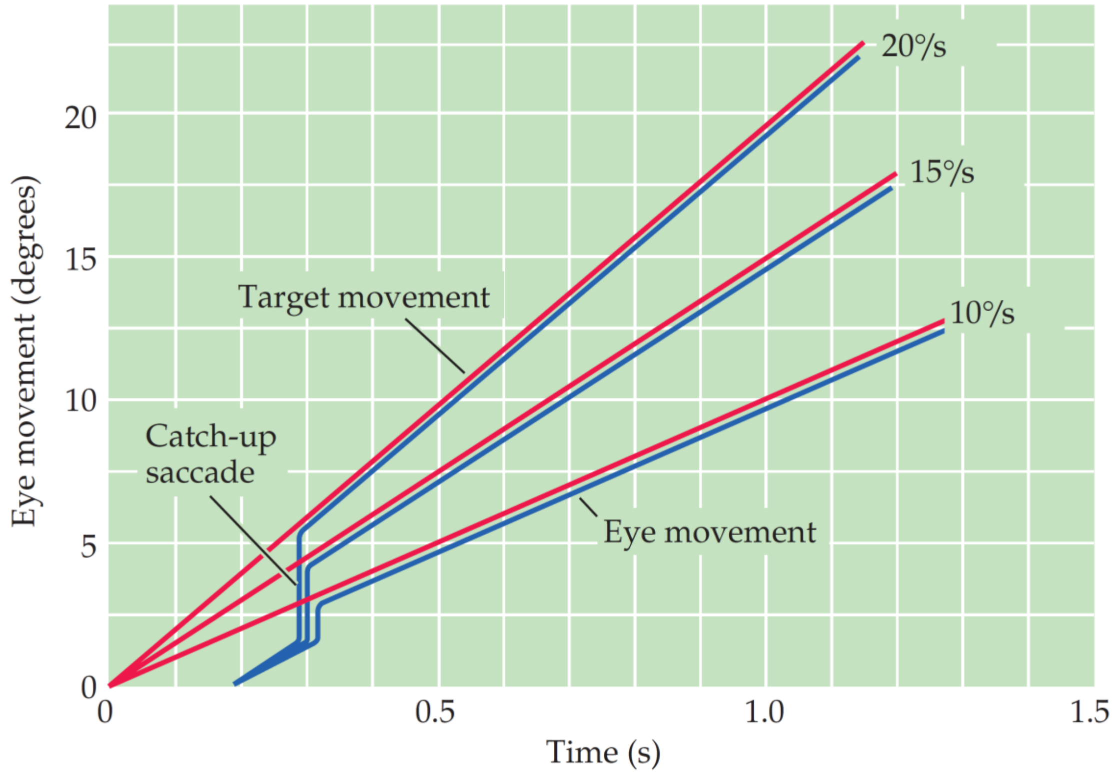

- Slower tracking movements that keep stimulus on the fovea
- Voluntary in that observer can choose whether or not to track moving stimulus
- Only highly trained observers can make a smooth pursuit movement in the absence of a moving target

Motion During Fixation

- **Drifts:**
  slow movements away from fixation point, 20 to 40 Hz
- **Flicks (microsaccades):**
  reposition the eye on target, 1 degree max
- **Ocular micro tremors:**
  150-2500nm, 40-100Hz

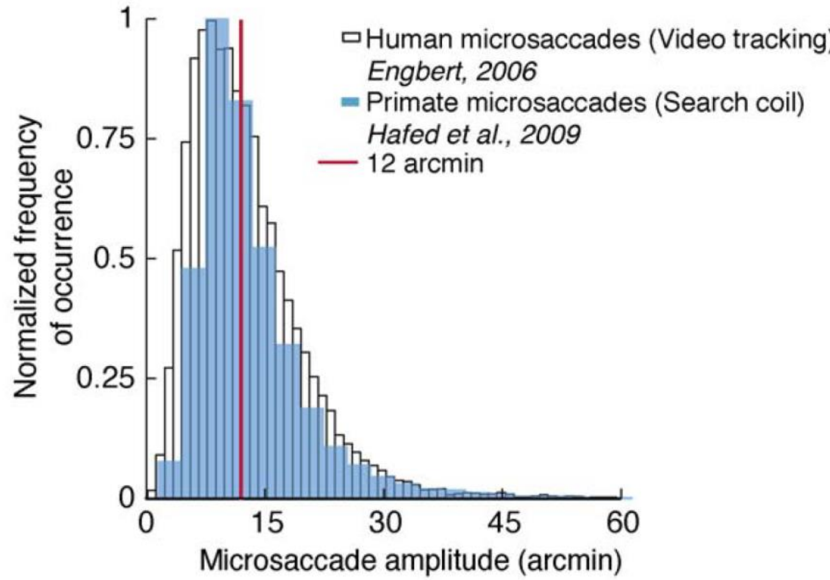

Cognitive Load Overview

From the Perspective of Computer Vision

** Each of the following bullet points have several papers validating it.*

- Pupil equations:
  - Brighter **light** = smaller **pupil**
  - Higher **cognitive load** = larger **pupil**
- Blink equations
  - Higher **cognitive load** = slower **blink rate**
  - Higher **cognitive load** = shorter **blink duration**

- **Questions:**
  - Which of these metrics can be accurately extracted in real-world driving data?
  - Are there other metrics that may work better in such conditions?

3D Convolutional Neural Networks

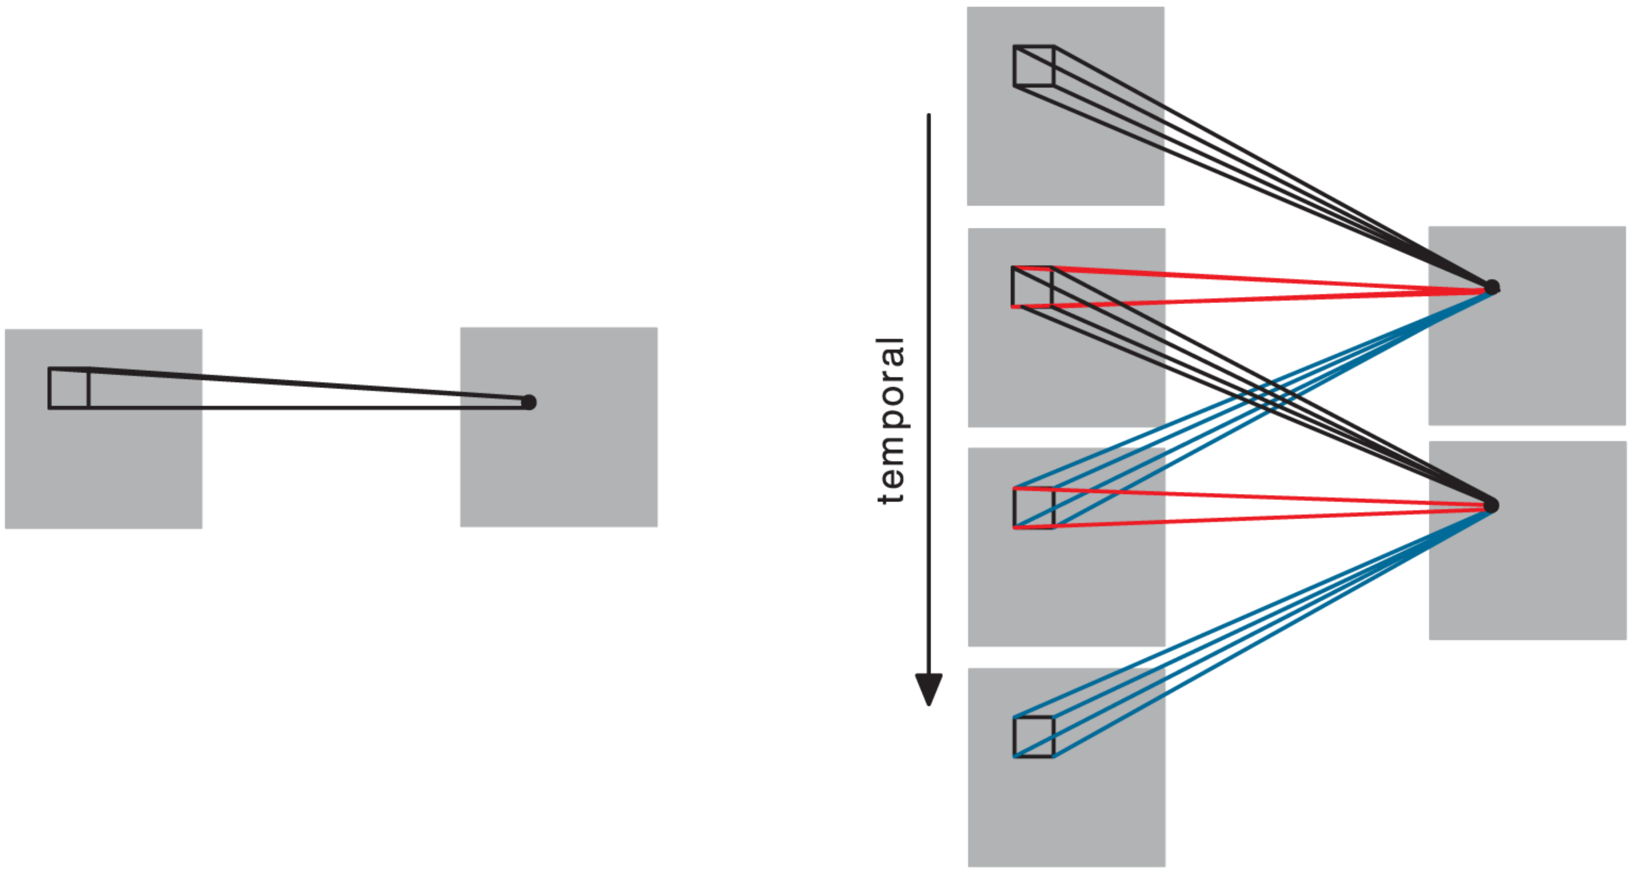

Real-World Data

92 drivers perform "n-back" tasks requiring various levels of cognitive load:

- **0-back:** Say the number right after it's read
- **1-back:** Say the number previous to the current one.
- **2-back:** Say the number 2 prior to the current one.

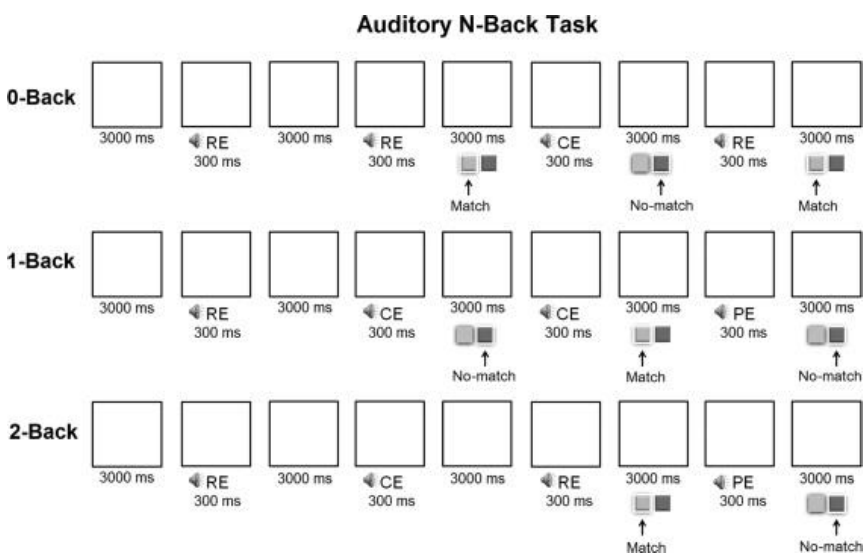

Cognitive Load Estimation

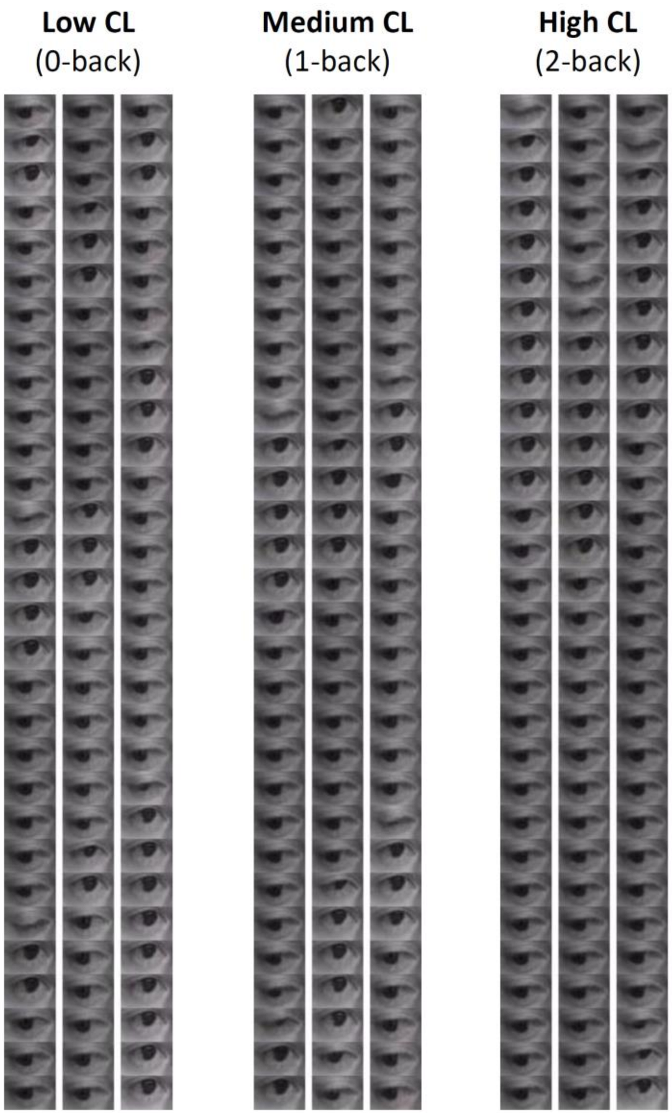

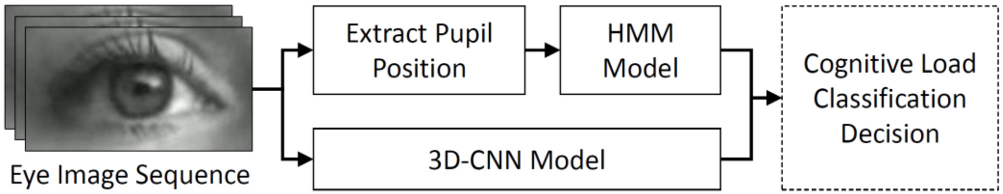

- 6 seconds, 16 fps, 90 images
- Two approaches: HMM and 3D-CNN

- **HMM:** Hidden Markov Model
  - **Input:** Sequence of pupil positions (normalized by intraocular segment)

- **3D-CNN:** Three Dimensional Convolutional Neural Network
  - **Input:** Sequence of raw images of eye region

Dealing with Vibration and Movement

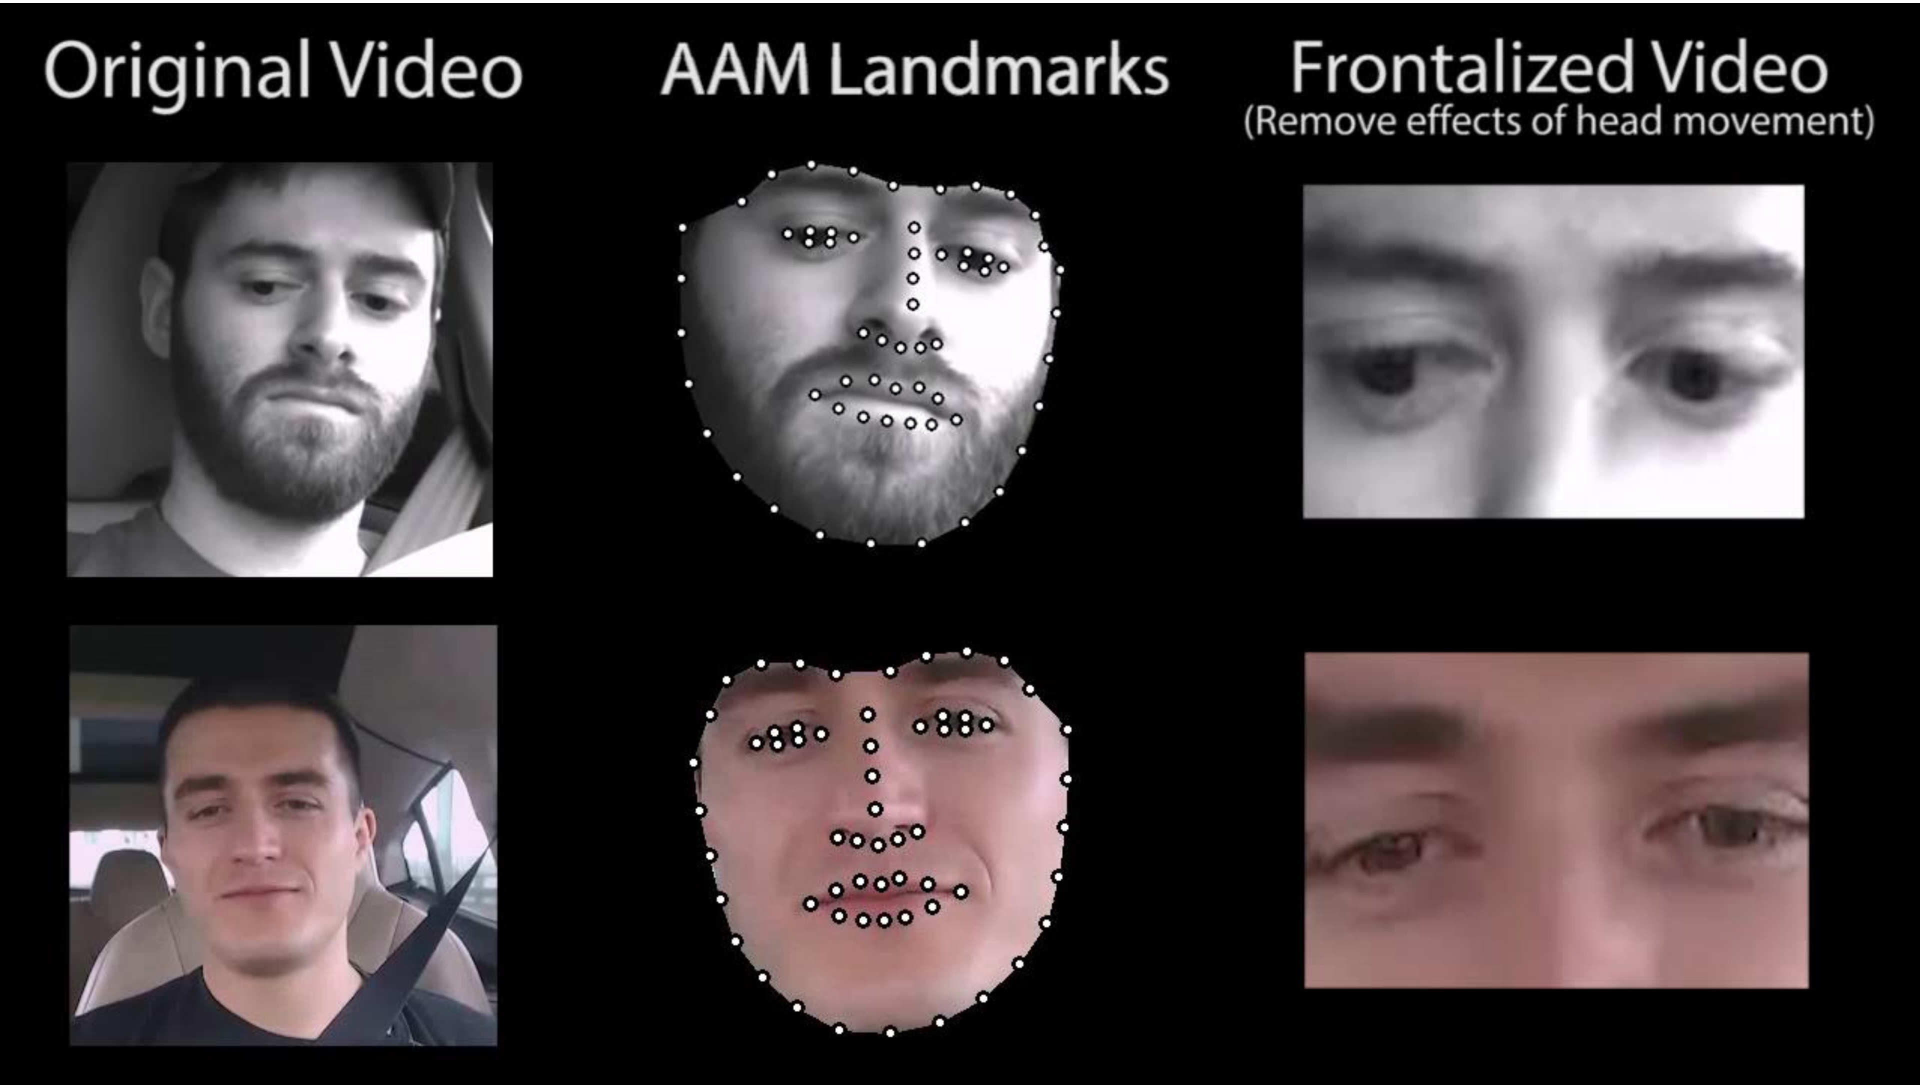

Preprocessing Pipeline

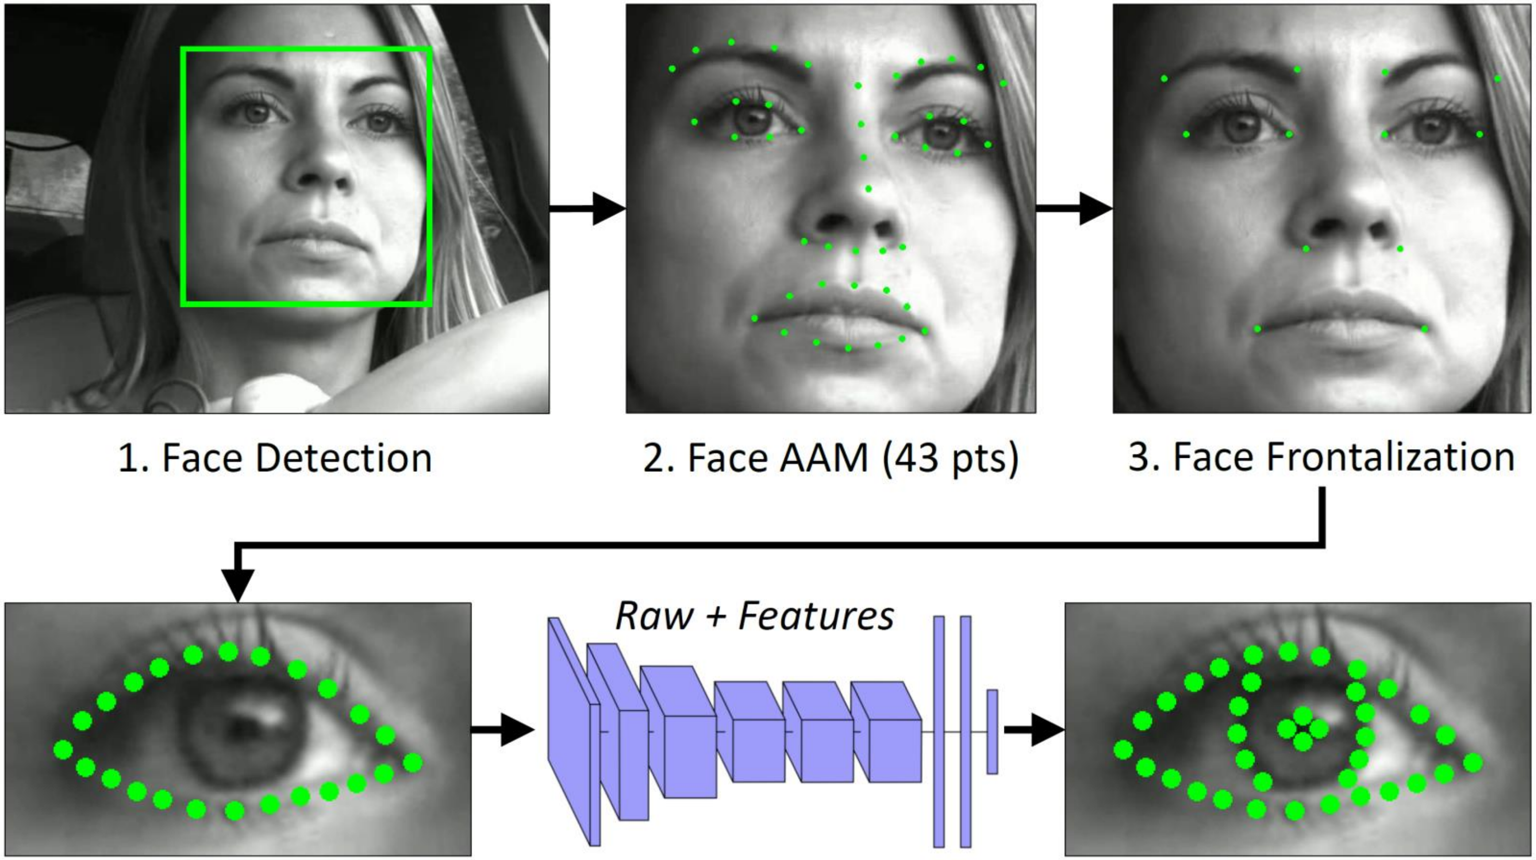

1. Face Detection
2. Face AAM (43 pts)
3. Face Frontalization
4. Eye Lid AAM (25 pts)
5. Classify Pupil Visibility
6. Pupil AAM (39 pts)

Visualizing the Dataset: Pupil Movement

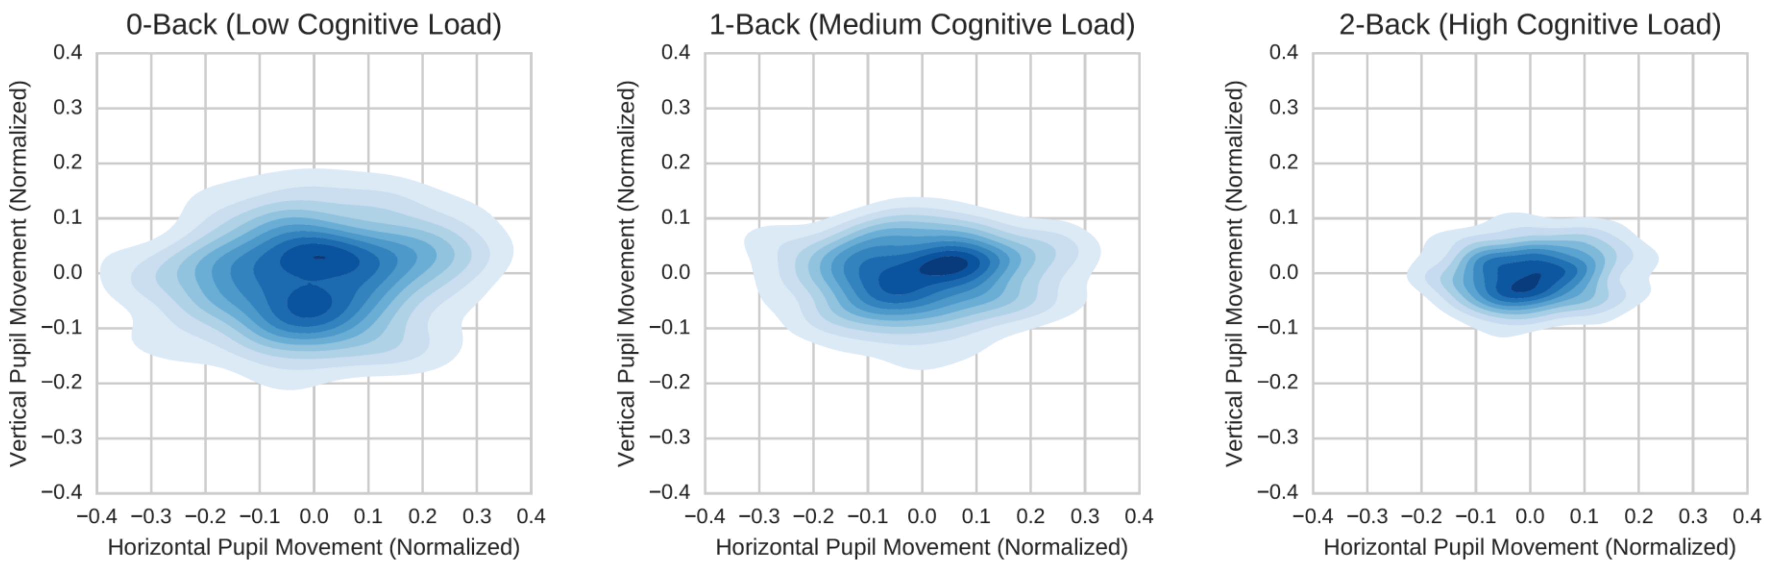

- **Metric:** Pupil position normalized by intraocular distance
- **Visualization:** Kernel density estimation (KDE)
- **Dataset size:** 92 subjects
- **Takeaway:** Observable aggregate differences between all 3 levels

Cognitive Load Estimation

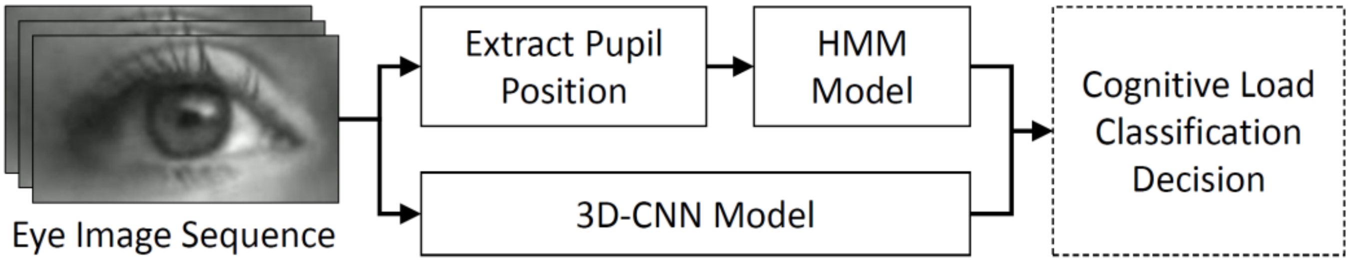

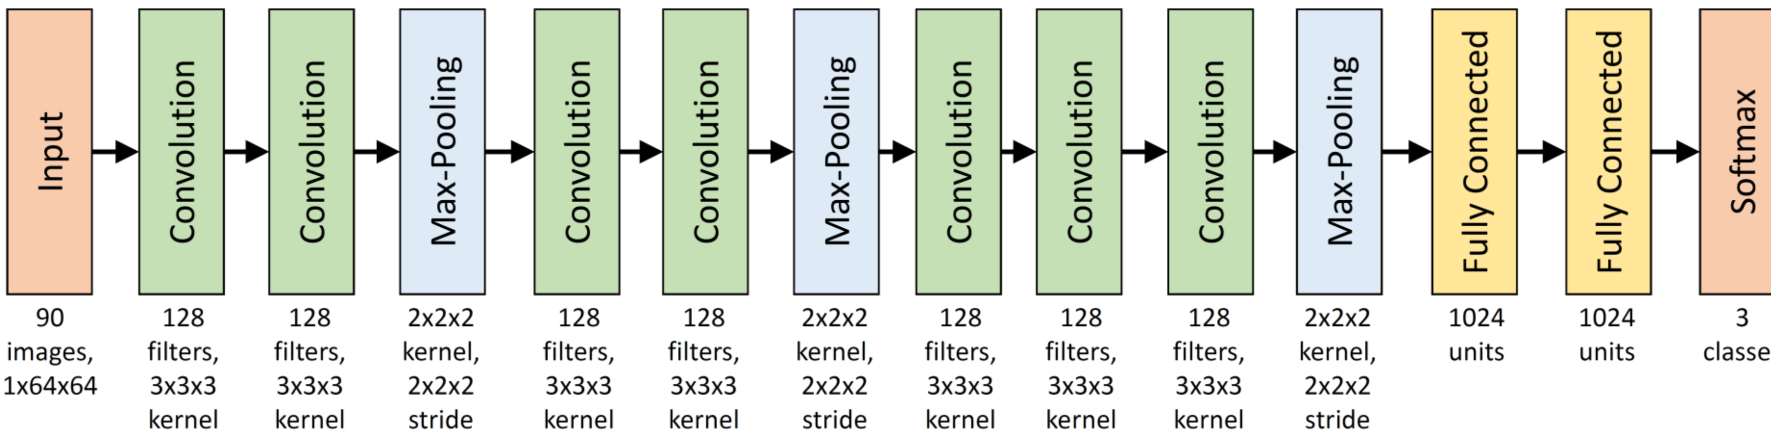

HMM: Hidden Markov Model

Input: Sequence of pupil positions (normalized by intraocular distance)

3D-CNN: Three Dimensional Convolutional Neural Network

Input: Sequence of raw images of eye region

Driver Cognitive Load Estimation

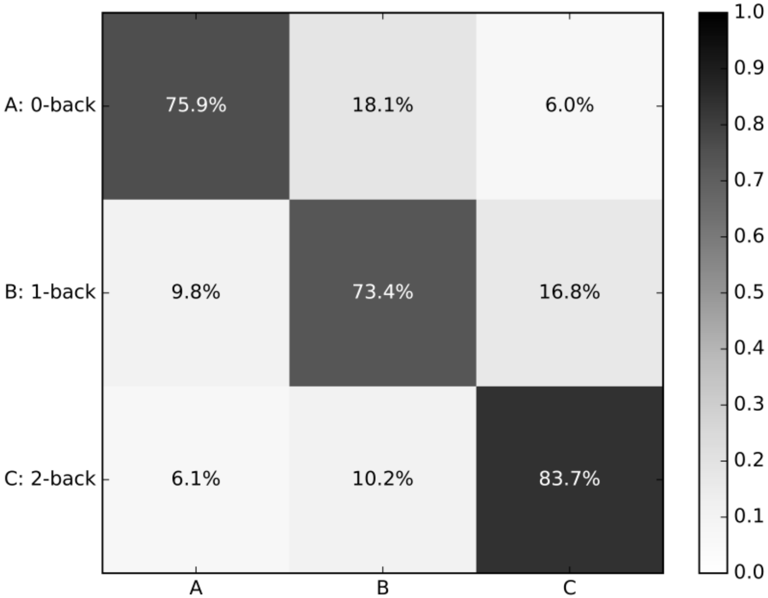

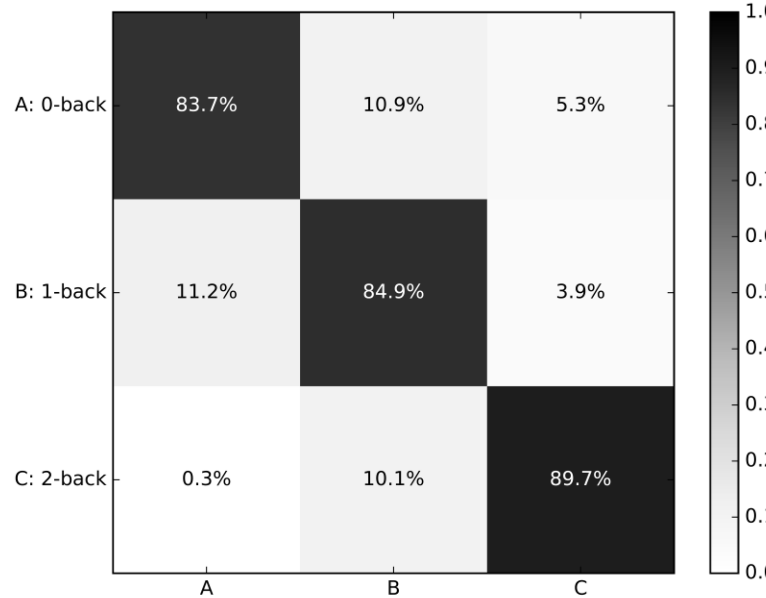

HMM Approach
Average Accuracy: **77.7%**

3D-CNN Approach
Average Accuracy: **86.1%**

Cognitive Load Estimation:
Open Source = Open Innovation

Implication: Make driver cognitive load estimation accessible

Webcam

Video Stream

Low
Cognitive Load

Medium
Cognitive Load

High
Cognitive Load

Real-Time Cognitive Load Estimation

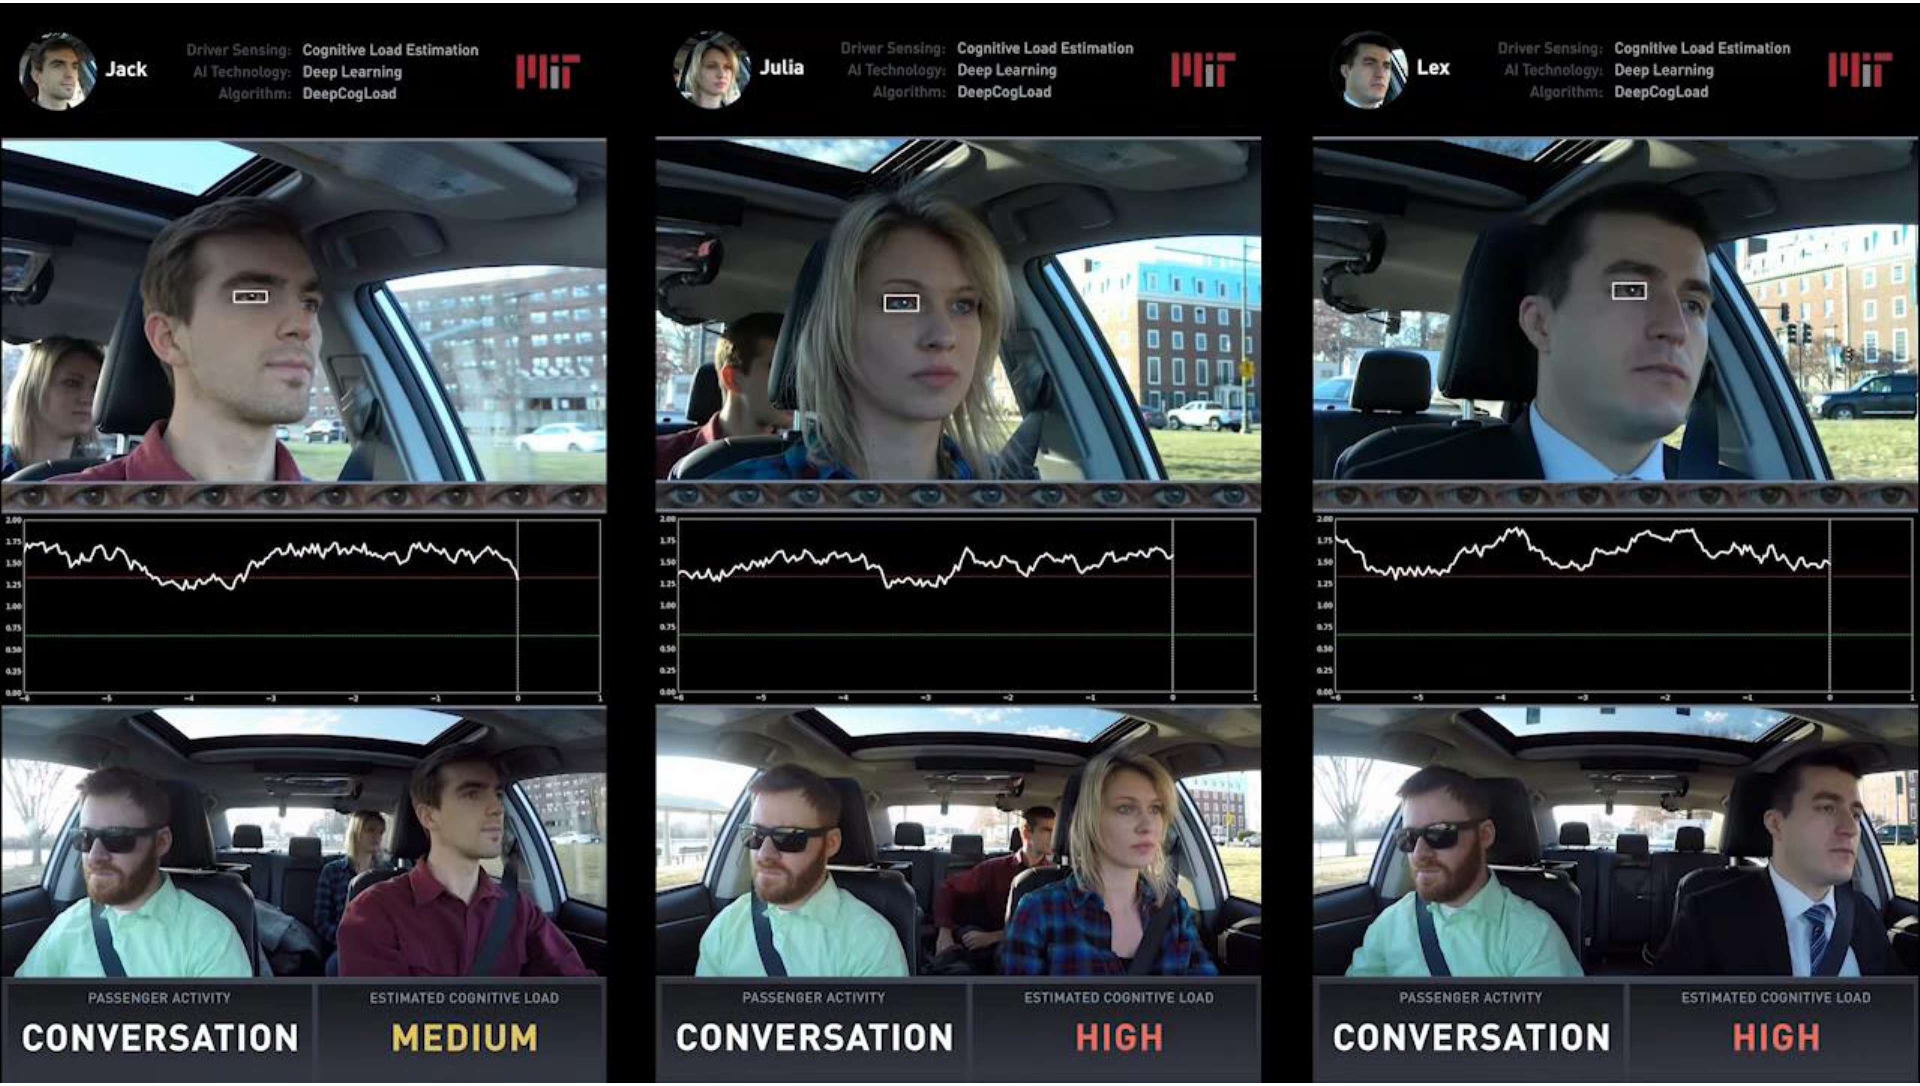

Overview

- Human Imperfections
- Pedestrian Detection
- Body Pose Estimation
- Face Detection
- Glance Classification
- Emotion Recognition
- Cognitive Load Estimation
- **Human-Centered Vision for Autonomous Vehicles**

Human-Centered Artificial Intelligence Approach

90% ← No — Human Needed — Yes → 10%

Solve the perception-control problem where **possible**:

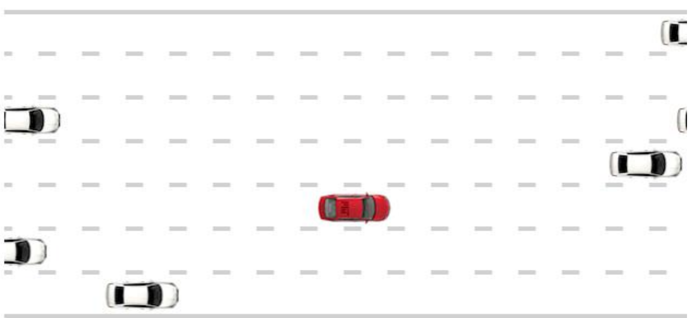

And where **not possible**: involve the human

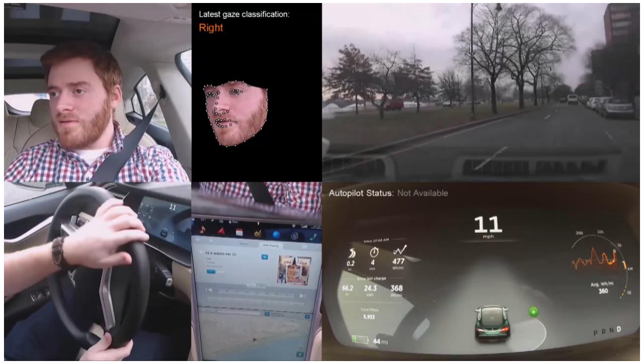

Human at the Center of Automation:
The Way to Full Autonomy Includes the Human

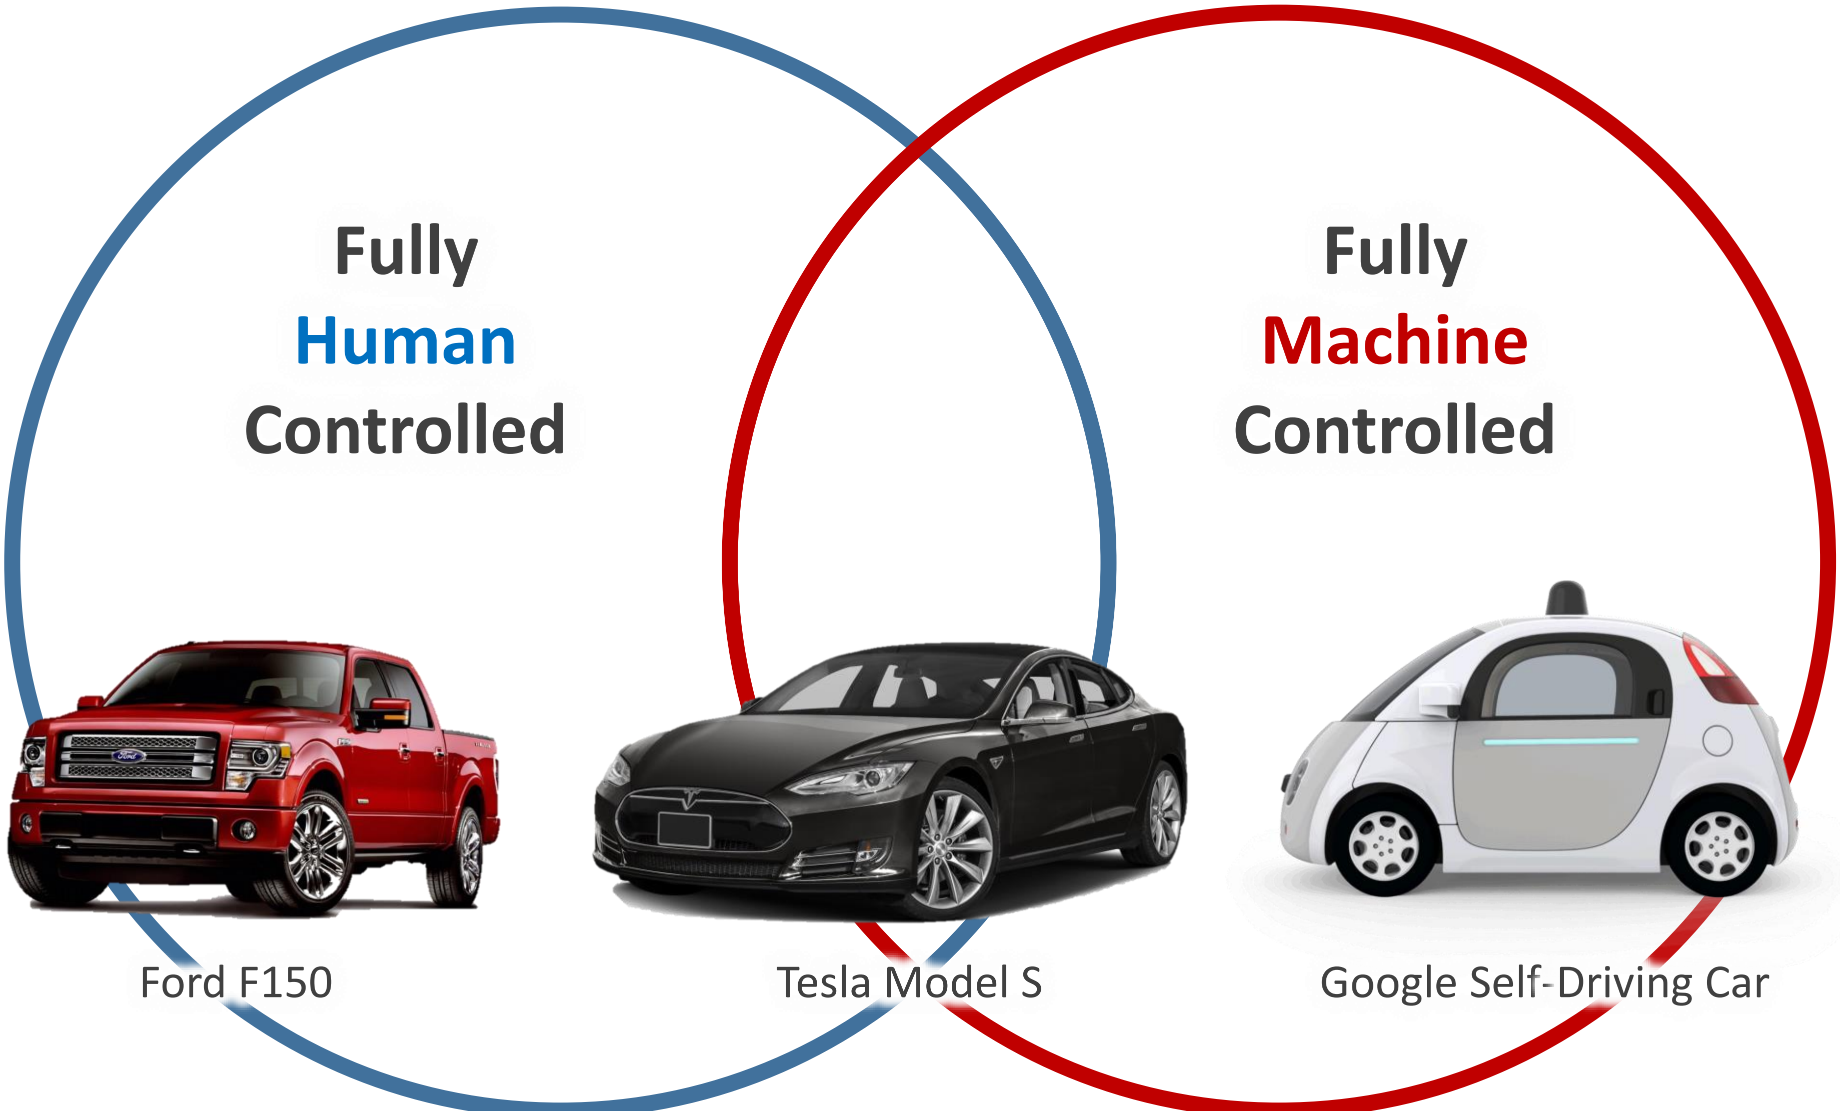

Stairway to Mass-Scale Automation

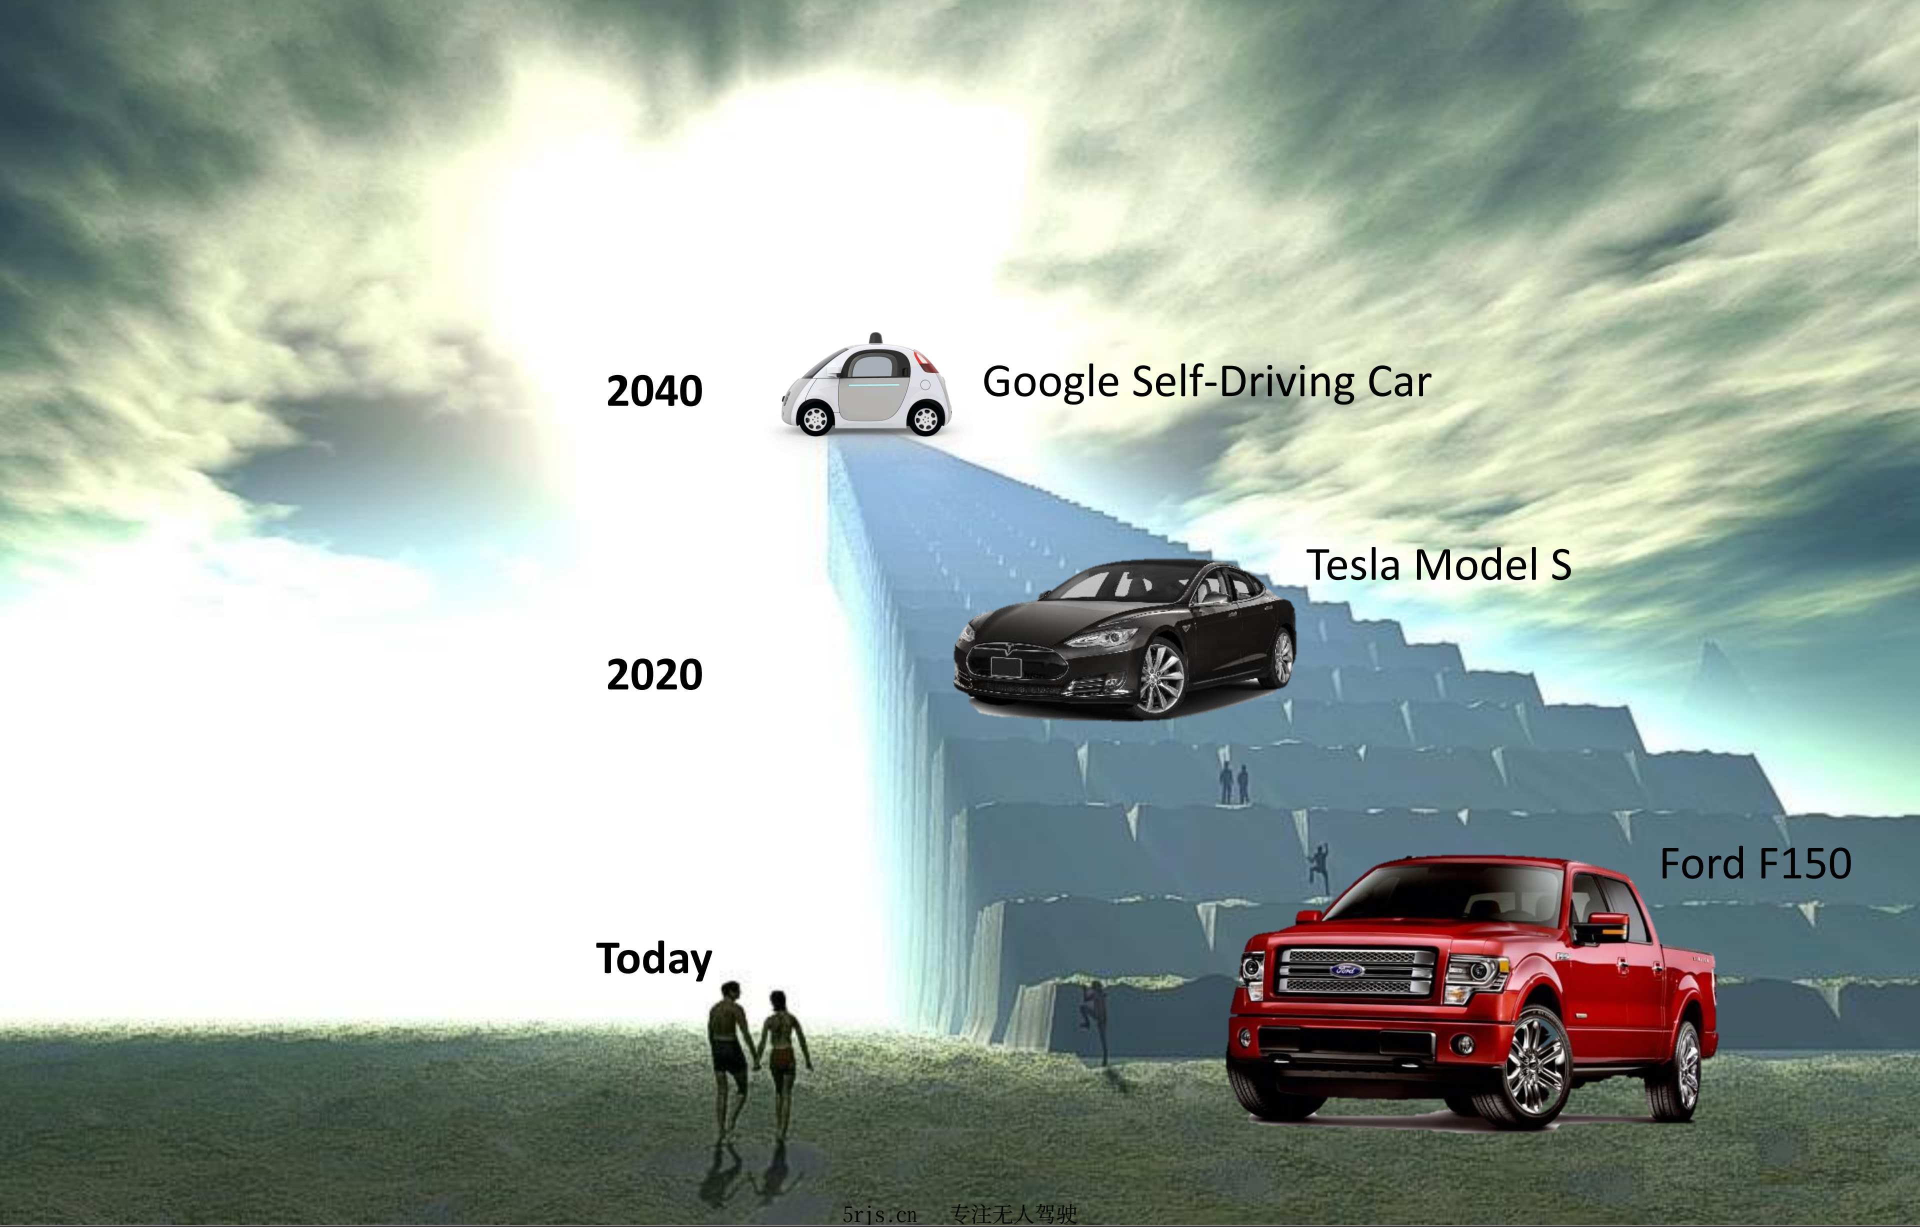

Stairway to Mass-Scale Automation

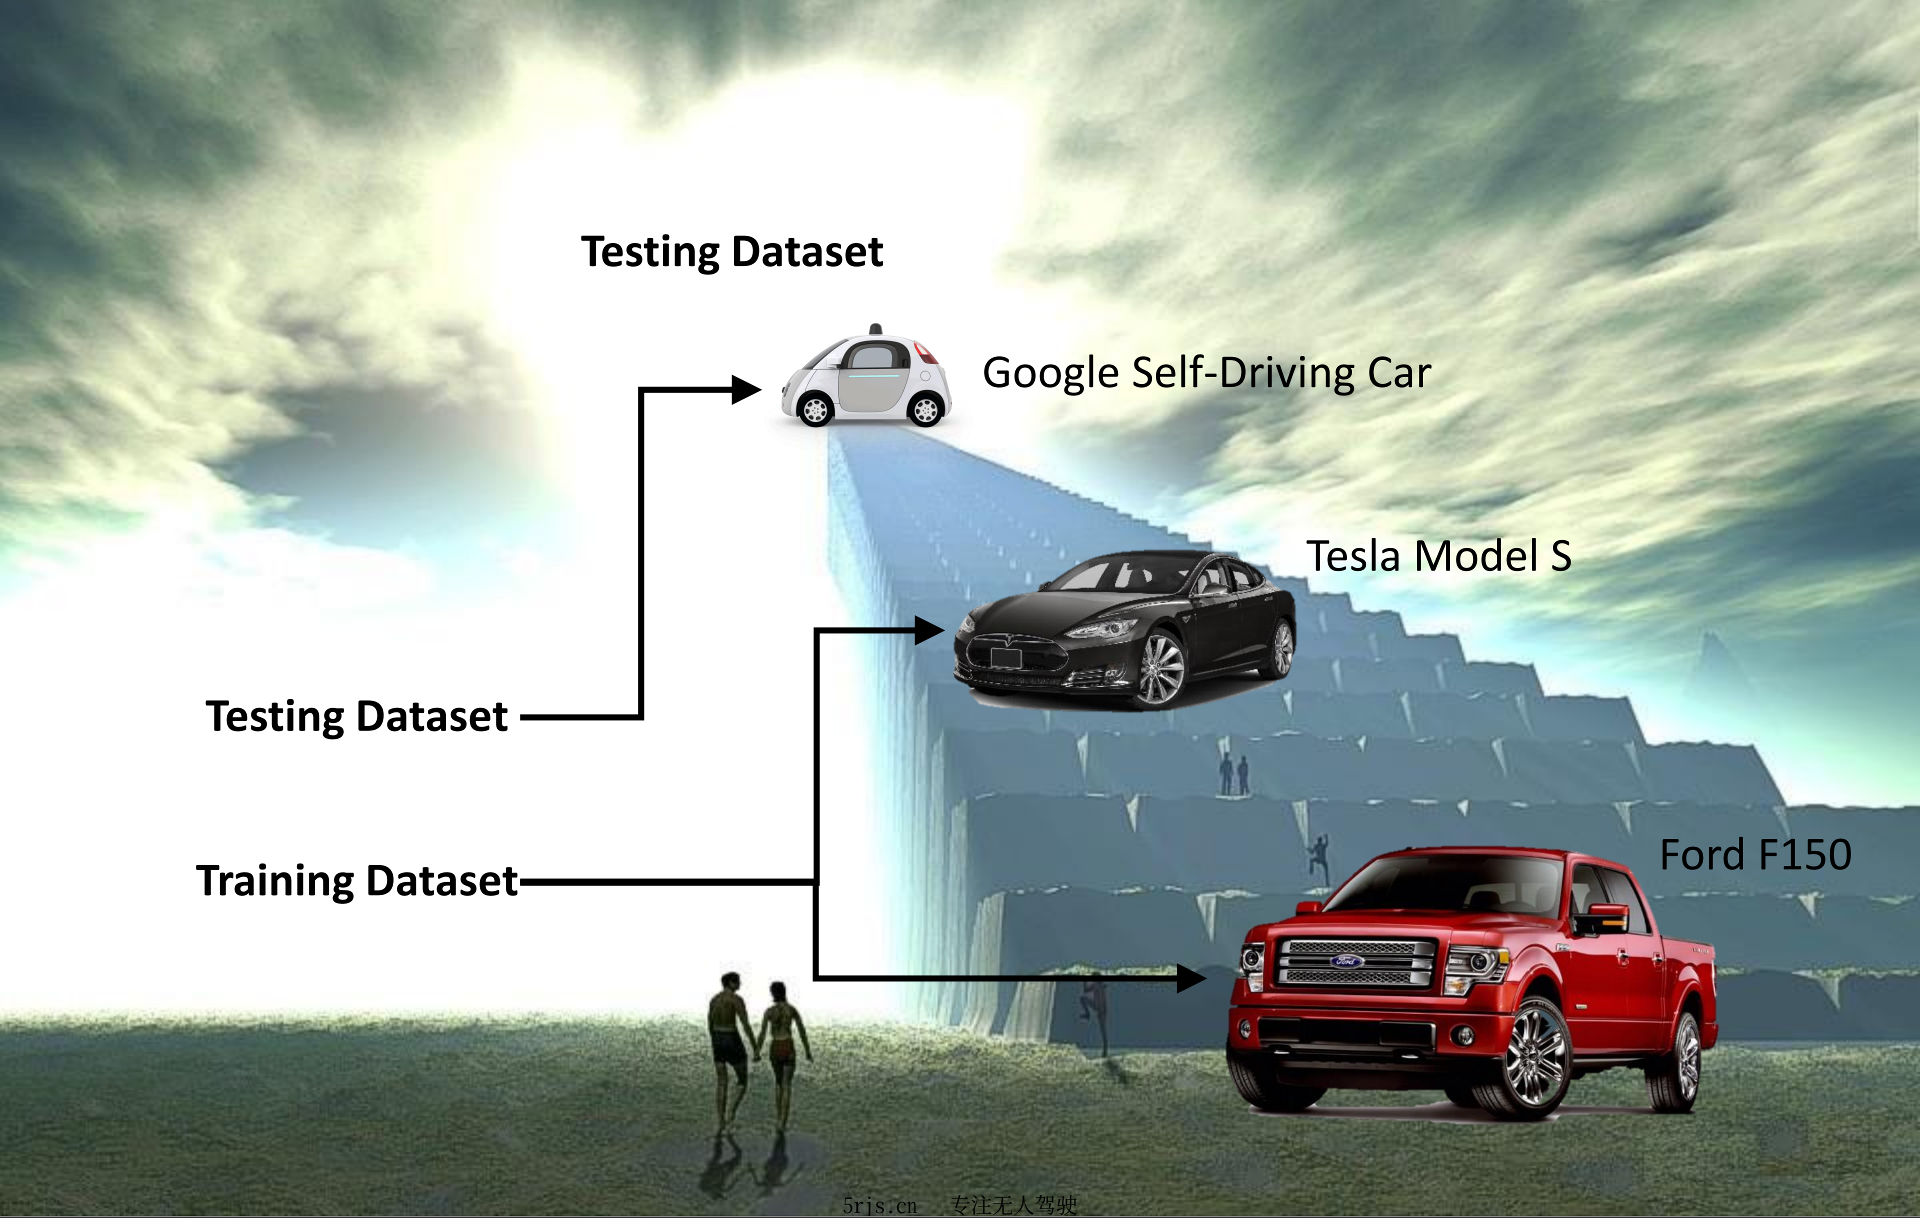

Human-Centered Autonomy

- A self-driving car may be more a Personal Robot and less a perfect Perception-Control system. Why:
  - **Flaws need humans:** The scene understanding problem requires much more than pixel-level labeling
  - **Exist with humans:** Achieving both an enjoyable and safe driving experience may require "driving like a human".

- Quite possibly, the first wide reaching and profound integration of personal robots in society.
  - **Wide reaching:** 1 billion cars on the road.
  - **Profound:** Human gives control of his/her life directly to robot.
  - **Personal:** One-on-one relationship of communication, collaboration, understanding and trust.

Human (and Machine) Imperfections

- “People call these things imperfections, but they’re not. That's the good stuff…”
- “And then we get to choose who we let in to our weird little worlds. You're not perfect, sport. And let me save you the suspense. This **girl** you met, **she** isn't perfect either. But the question is: whether or not you're perfect for each other. That's the whole deal. That's what intimacy is all about…”
- “Now you can know everything in the world, sport, but the only way you're finding out that one is by giving it a shot.”

MIT HCAV: Human-Centered Autonomous Vehicle

March 2018

CHI 2018 Course:
Deep Learning for Understanding the Human

CHI 2018
Engage with CHI

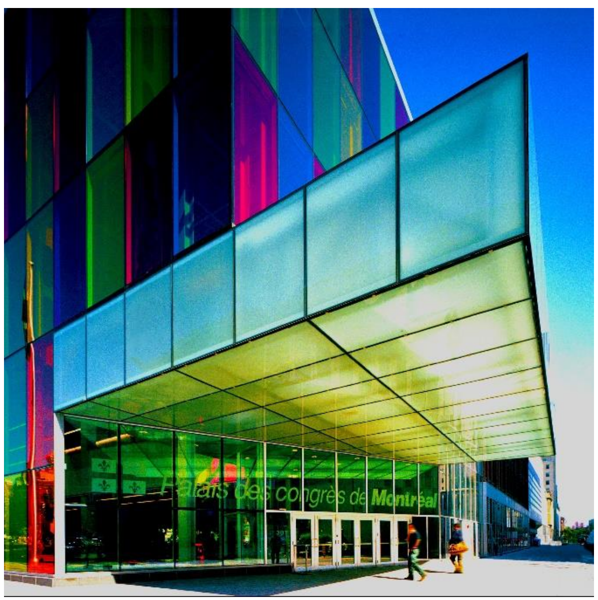

- Part 1 (80 minutes)
  - Introduction to Deep Learning
    - Theory, insights, and intuitions
    - Tools to get started applying DL to various domains
  - Convolutional Neural Networks
    - Face recognition
    - Eye tracking
    - Cognitive load estimation
    - Emotion recognition
- Part 2 (80 minutes)
  - Recurrent Neural Networks
    - Natural Language Processing
    - Voice Recognition
  - Mixing Convolutional and Recurrent Neural Networks
    - Activity recognition
- Part 3 (80 minutes)
  - Generative Neural Networks
    - Speech Synthesis
    - Peripheral Vision Visualization

* dates, times, rooms in red are different than the usual

6.S099
Artificial General Intelligence
agi.mit.edu

Mon, Jan 22 7pm, 54-100	Lex Fridman, MIT	Artificial General Intelligence
Tue, Jan 23 7pm, 54-100	Josh Tenenbaum, MIT	Computational Cognitive Science
Wed, Jan 24 1pm, 10-250	Ray Kurzweil, Google	How to Create a Mind
Thu, Jan 25 7pm, 54-100	Lisa Feldman Barrett, NEU	Emotion Creation
Fri, Jan 26 7pm, 54-100	Nate Derbinsky, NEU	Cognitive Modeling
Mon, Jan 29 1:30pm, 26-100	Andrej Karpathy, Tesla	Deep Learning
Mon, Jan 29 7pm, 54-100	Stephen Wolfram, Wolfram Research	Knowledge-Based Programming
Tue, Jan 30 7pm, 54-100	Richard Moyes, Article36	AI Safety: Autonomous Weapon Systems
Wed, Jan 31 7pm, 54-100	Marc Raibert, Boston Dynamics	Robots That Work in the Real World
Thu, Feb 1 7pm, 54-100	Ilya Sutskever, OpenAI	Deep Reinforcement Learning
Fri, Feb 2 7pm, 54-100	Lex Fridman, MIT	Human-Centered Artificial Intelligence

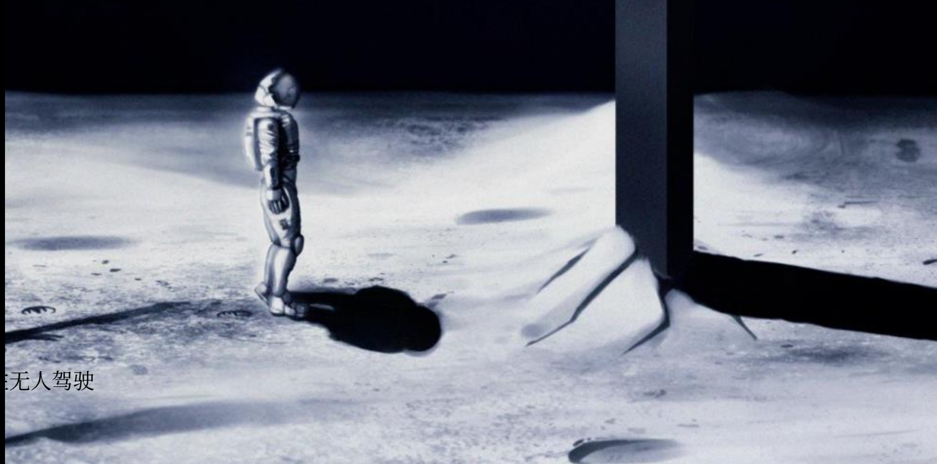

无人驾驶

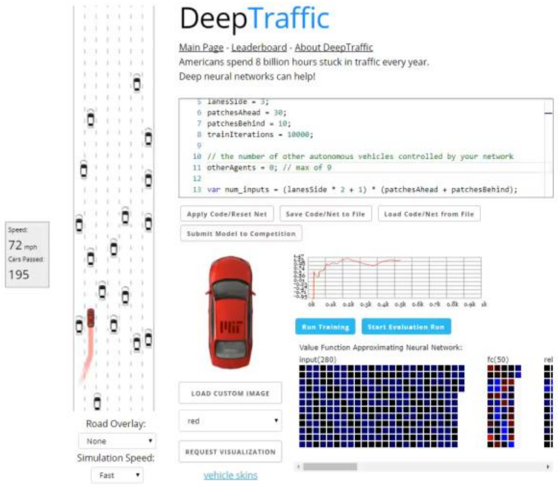

What Next?

- Competitions
  - Ongoing until May 2018. Results, insights → NIPS 2018
  - DeepTraffic: https://selfdrivingcars.mit.edu/deeptraffic
  - SegFuse: https://selfdrivingcars.mit.edu/segfuse
  - DeepCrash: https://selfdrivingcars.mit.edu/deepcrash

- Upcoming MIT Courses:
  - 6.S099: Artificial General Intelligence https://agi.mit.edu
  - 6.S191: Introduction to Deep Learning: http://introtodeeplearning.com
  - 15.S14: Global Business of AI & Robotics http://tiny.cc/gbair18

- If you're interested in the application of deep learning in the automotive space, come do research with us: https://hcai.mit.edu/join *(opens in Feb 2018)*

Thank You

NVIDIA

Google

amazon alexa

auto

Autoliv

CSRC
Collaborative Safety Research Center
TOYOTA

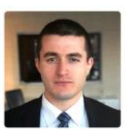

Lex Fridman
Instructor

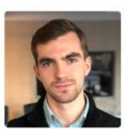

Jack Terwilliger
TA

Julia Kindelsberger
TA

Dan Brown
TA

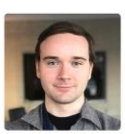

Michael Glazer
TA

Li Ding
TA

Spencer Dodd
TA

Benedikt Jenik
TA